



Junos[®] OS

Network Management Administration Guide



Modified: 2017-11-07

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. and/or its affiliates in the United States and other countries. All other trademarks may be 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.

Junos® OS Network Management Administration Guide
Copyright © 2017 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

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 <http://www.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 the Documentation	xxxi
	Documentation and Release Notes	xxxi
	Supported Platforms	xxxi
	Using the Examples in This Manual	xxxii
	Merging a Full Example	xxxii
	Merging a Snippet	xxxiii
	Documentation Conventions	xxxiii
	Documentation Feedback	xxxv
	Requesting Technical Support	xxxvi
	Self-Help Online Tools and Resources	xxxvi
	Opening a Case with JTAC	xxxvi
Part 1	Overview	
Chapter 1	Network Management Overview	3
	Understanding Device Management Functions in Junos OS	3
	Understanding the Integrated Local Management Interface	6
Chapter 2	Introduction to Network Monitoring	7
	Monitoring Overview	7
	Diagnostic Tools Overview	8
	J-Web Diagnostic Tools	9
	CLI Diagnostic Commands	9
Part 2	Network Monitoring Using SNMP	
Chapter 3	SNMP Overview	15
	Understanding SNMP Implementation in Junos OS	15
	SNMPv3 Overview	18
Chapter 4	SNMP MIBs and Traps Supported by Junos OS	21
	Enterprise-Specific SNMP MIBs Supported by Junos OS	21
	Standard SNMP MIBs Supported by Junos OS	32
	Standard SNMP Traps Supported by Junos OS	50
	Standard SNMP Version 1 Traps	50
	Standard SNMP Version 2 Traps	53
	Enterprise-Specific SNMP Traps Supported by Junos OS	57
	Juniper Networks Enterprise-Specific SNMP Version 1 Traps	57
	Juniper Networks Enterprise-Specific SNMP Version 2 Traps	63
Chapter 5	Loading MIB Files to a Network Management System	73
	Loading MIB Files to a Network Management System	73

Chapter 6	Configuring SNMP	77
	Configuration Statements at the [edit snmp] Hierarchy Level	78
	Optimizing the Network Management System Configuration for the Best Results	81
	Changing the Polling Method from Column-by-Column to Row-by-Row	82
	Reducing the Number of Variable Bindings per PDU	82
	Increasing Timeout Values in Polling and Discovery Intervals	82
	Reducing Incoming Packet Rate at the snmpd	82
	Configuring Options on Managed Devices for Better SNMP Response Time	83
	Enabling the stats-cache-lifetime Option	83
	Filtering Out Duplicate SNMP Requests	83
	Excluding Interfaces That Are Slow in Responding to SNMP Queries	84
	Configuring SNMP on Devices Running Junos OS	85
	Configuring Basic Settings for SNMPv1 and SNMPv2	85
	Configuring Basic Settings for SNMPv3	85
	Configuring System Name, Location, Description, and Contact Information	87
	Configuring the System Contact on a Device Running Junos OS	88
	Configuring the System Location for a Device Running Junos OS	89
	Configuring the System Description on a Device Running Junos OS	89
	Configuring SNMP Details	90
	Configuring a Different System Name	91
	Configuring the Commit Delay Timer	92
	Filtering Duplicate SNMP Requests	93
	Configuring SNMP Communities	93
	Examples: Configuring the SNMP Community String	96
	Adding a Group of Clients to an SNMP Community	97
	Configuring a Proxy SNMP Agent	98
	Configuring SNMP Traps	100
	Configuring SNMP Trap Options and Groups on a Device Running Junos OS	102
	Configuring SNMP Trap Options	103
	Configuring the Source Address for SNMP Traps	103
	Configuring the Agent Address for SNMP Traps	105
	Adding snmpTrapEnterprise Object Identifier to Standard SNMP Traps	106
	Configuring SNMP Trap Groups	106
	Example: Configuring SNMP Trap Groups	109
	Configuring the Interfaces on Which SNMP Requests Can Be Accepted	109
	Example: Configuring Secured Access List Checking	110
	Filtering Interface Information Out of SNMP Get and GetNext Output	110
	Configuring MIB Views	111
	Configuring Ping Proxy MIB	113
Chapter 7	Configuring SNMPv3	115
	Minimum SNMPv3 Configuration on a Device Running Junos OS	116
	Example: SNMPv3 Configuration	117
	Configuring the Local Engine ID	120

Creating SNMPv3 Users	121
Example: Creating SNMPv3 Users	122
Configuring the SNMPv3 Authentication Type	123
Configuring MD5 Authentication	123
Configuring SHA Authentication	123
Configuring No Authentication	124
Configuring the SNMPv3 Encryption Type	124
Configuring the Advanced Encryption Standard Algorithm	124
Configuring the Data Encryption Algorithm	125
Configuring Triple DES	125
Configuring No Encryption	126
Defining Access Privileges for an SNMP Group	126
Configuring the Access Privileges Granted to a Group	127
Configuring the Group	127
Configuring the Security Model	128
Configuring the Security Level	128
Associating MIB Views with an SNMP User Group	129
Configuring the Notify View	129
Configuring the Read View	130
Configuring the Write View	130
Example: Configuring the Access Privileges Granted to a Group	130
Assigning Security Model and Security Name to a Group	131
Configuring the Security Model	132
Assigning Security Names to Groups	132
Configuring the Group	132
Example: Security Group Configuration	133
Configuring SNMPv3 Traps on a Device Running Junos OS	133
Configuring the SNMPv3 Trap Notification	135
Example: Configuring SNMPv3 Trap Notification	136
Configuring the Trap Notification Filter	136
Configuring the Trap Target Address	137
Configuring the Address	138
Configuring the Address Mask	138
Configuring the Port	138
Configuring the Routing Instance	138
Configuring the Trap Target Address	139
Applying Target Parameters	139
Example: Configuring the Tag List	140
Defining and Configuring the Trap Target Parameters	140
Applying the Trap Notification Filter	141
Configuring the Target Parameters	141
Configuring the Message Processing Model	142
Configuring the Security Model	142
Configuring the Security Level	142
Configuring the Security Name	143
Configuring SNMP Informs	143
Configuring the Remote Engine and Remote User	144
Example: Configuring the Remote Engine ID and Remote User	145
Configuring the Inform Notification Type and Target Address	148

	Example: Configuring the Inform Notification Type and Target Address	150
	Configuring the SNMPv3 Community	151
	Configuring the Community Name	151
	Configuring the Context	152
	Configuring the Security Names	152
	Configuring the Tag	152
	Example: Configuring an SNMPv3 Community	153
Chapter 8	Configuring SNMP for Routing Instances	155
	Understanding SNMP Support for Routing Instances	155
	SNMP MIBs Supported for Routing Instances	156
	Support Classes for MIB Objects	166
	SNMP Traps Supported for Routing Instances	167
	Identifying a Routing Instance	168
	Enabling SNMP Access over Routing Instances	169
	Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community	169
	Example: Configuring Interface Settings for a Routing Instance	170
	Configuring Access Lists for SNMP Access over Routing Instances	172
Chapter 9	Configuring SNMP Remote Operations	173
	SNMP Remote Operations Overview	173
	SNMP Remote Operation Requirements	174
	Setting SNMP Views	174
	Example: Setting SNMP Views	174
	Setting Trap Notification for Remote Operations	175
	Example: Setting Trap Notification for Remote Operations	175
	Using Variable-Length String Indexes	175
	Example: Set Variable-Length String Indexes	175
	Enabling Logging	176
	Using the Ping MIB for Remote Monitoring Devices Running Junos OS	176
	Starting a Ping Test	177
	Using Multiple Set Protocol Data Units (PDUs)	177
	Using a Single Set PDU	177
	Monitoring a Running Ping Test	178
	pingResultsTable	178
	pingProbeHistoryTable	179
	Generating Traps	180
	Gathering Ping Test Results	181
	Stopping a Ping Test	182
	Interpreting Ping Variables	182
	Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS	183
	Starting a Traceroute Test	184
	Using Multiple Set PDUs	184
	Using a Single Set PDU	184
	Monitoring a Running Traceroute Test	185
	traceRouteResultsTable	185
	traceRouteProbeResultsTable	186
	traceRouteHopsTable	187

	Generating Traps	188
	Monitoring Traceroute Test Completion	189
	Gathering Traceroute Test Results	190
	Stopping a Traceroute Test	191
	Interpreting Traceroute Variables	192
Chapter 10	Tracing SNMP Activity	195
	Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance	
	on a Device Running Junos OS	195
	Checking for MIB Objects Registered with the snmpd	195
	Tracking SNMP Activity	197
	Monitoring SNMP Statistics	199
	Checking CPU Utilization	199
	Checking Kernel and Packet Forwarding Engine Response	200
	Tracing SNMP Activity on a Device Running Junos OS	201
	Configuring the Number and Size of SNMP Log Files	202
	Configuring Access to the Log File	203
	Configuring a Regular Expression for Lines to Be Logged	203
	Configuring the Trace Operations	203
	Example: Tracing SNMP Activity	205
Chapter 11	SNMP FAQs	207
	Junos OS SNMP FAQ Overview	207
	Junos OS SNMP FAQs	208
	Junos OS SNMP Support FAQs	208
	Junos OS MIBs FAQs	209
	Junos OS SNMP Configuration FAQs	216
	SNMPv3 FAQs	220
	SNMP Interaction with Juniper Networks Devices FAQs	222
	SNMP Traps and Informs FAQs	224
	Junos OS Dual Routing Engine Configuration FAQs	230
	SNMP Support for Routing Instances FAQs	231
	SNMP Counters FAQs	232
Part 3	Remote Monitoring (RMON) with SNMP	
Chapter 12	RMON Overview	237
	Understanding RMON Alarms	237
	alarmTable	237
	jnxRmonAlarmTable	238
	Understanding RMON Events	238
	eventTable	239
Chapter 13	Configuring RMON Alarms and Events	241
	Understanding RMON Alarms and Events Configuration	241
	Minimum RMON Alarm and Event Entry Configuration	242
	Configuring an RMON Alarm Entry and Its Attributes	242
	Configuring the Alarm Entry	243
	Configuring the Description	243
	Configuring the Falling Event Index or Rising Event Index	243

	Configuring the Falling Threshold or Rising Threshold	244
	Configuring the Interval	244
	Configuring the Falling Threshold Interval	244
	Configuring the Request Type	245
	Configuring the Sample Type	245
	Configuring the Startup Alarm	245
	Configuring the System Log Tag	246
	Configuring the Variable	246
	Configuring an RMON Event Entry and Its Attributes	246
	Example: Configuring an RMON Alarm and Event Entry	247
Chapter 14	Monitoring RMON Alarms and Events	249
	Using alarmTable to Monitor MIB Objects	249
	Creating an Alarm Entry	249
	Configuring the Alarm MIB Objects	249
	alarmInterval	250
	alarmVariable	250
	alarmSampleType	250
	alarmValue	250
	alarmStartupAlarm	250
	alarmRisingThreshold	251
	alarmFallingThreshold	251
	alarmOwner	251
	alarmRisingEventIndex	251
	alarmFallingEventIndex	252
	Activating a New Row in alarmTable	252
	Modifying an Active Row in alarmTable	252
	Deactivating a Row in alarmTable	252
	Using eventTable to Log Alarms	252
	Creating an Event Entry	253
	Configuring the MIB Objects	253
	eventType	253
	eventCommunity	253
	eventOwner	254
	eventDescription	254
	Activating a New Row in eventTable	254
	Deactivating a Row in eventTable	255
Chapter 15	Using RMON to Monitor Network Service Quality	257
	Understanding RMON for Monitoring Service Quality	257
	Setting Thresholds	257
	RMON Command-Line Interface	258
	RMON Event Table	259
	RMON Alarm Table	259
	Troubleshooting RMON	260
	Understanding Measurement Points, Key Performance Indicators, and Baseline Values	261
	Measurement Points	261
	Basic Key Performance Indicators	262
	Setting Baselines	262

	Defining and Measuring Network Availability	262
	Defining Network Availability	263
	Monitoring the SLA and the Required Bandwidth	264
	Measuring Availability	265
	Real-Time Performance Monitoring	265
	Measuring Health	268
	Measuring Performance	274
	Measuring Class of Service	277
	Inbound Firewall Filter Counters per Class	278
	Monitoring Output Bytes per Queue	279
	Dropped Traffic	280
Part 4	Health Monitoring with SNMP	
Chapter 16	Configuring Health Monitoring	285
	Configuring Health Monitoring on Devices Running Junos OS	285
	Monitored Objects	286
	Minimum Health Monitoring Configuration	287
	Configuring the Falling Threshold or Rising Threshold	287
	Configuring the Interval	288
	Log Entries and Traps	288
	Example: Configuring Health Monitoring	288
Part 5	Gathering Statistics for Accounting Purposes Using Accounting Options, Source Class Usage and Destination Class Usage Options	
Chapter 17	Accounting Options, Source Class Usage and Destination Class Usage Options Overview	291
	Accounting Options Overview	291
	Understanding Source Class Usage and Destination Class Usage Options	292
Chapter 18	Configuring Accounting Options, Source Class Usage and Destination Class Usage Options	295
	Configuration Statements at the [edit accounting-options] Hierarchy Level	295
	Accounting Options Configuration	296
	Accounting Options—Full Configuration	297
	Minimum Accounting Options Configuration	300
	Configuring Accounting-Data Log Files	304
	Configuring How Long Backup Files Are Retained	305
	Configuring the Maximum Size of the File	305
	Configuring Archive Sites for the Files	306
	Configuring Local Backup for Accounting Files	306
	Configuring Files to Be Compressed	307
	Configuring the Maximum Number of Files	307
	Configuring the Storage Location of the File	307
	Configuring Files to Be Saved After a Change in Mastership	308
	Configuring the Start Time for File Transfer	308

Configuring the Transfer Interval of the File	308
Configuring the Interface Profile	309
Configuring Fields	310
Configuring the File Information	310
Configuring Cleared Statistics to be Reported in the Flat File	310
Configuring the Interval	311
Example: Configuring the Interface Profile	311
Configuring the Filter Profile	312
Configuring the Counters	313
Configuring the File Information	313
Configuring the Interval	314
Example: Configuring a Filter Profile	314
Example: Configuring Interface-Specific Firewall Counters and Filter Profiles ..	315
Configuring SCU or DCU	316
Creating Prefix Route Filters in a Policy Statement	317
Applying the Policy to the Forwarding Table	317
Enabling Accounting on Inbound and Outbound Interfaces	317
Configuring SCU on a Virtual Loopback Tunnel Interface	318
Example: Configuring a Virtual Loopback Tunnel Interface on a Provider Edge Router Equipped with a Tunnel PIC	319
Example: Mapping the VRF Instance Type to the Virtual Loopback Tunnel Interface	319
Example: Sending Traffic Received from the Virtual Loopback Interface Out the Source Class Output Interface	320
Configuring Class Usage Profiles	320
Configuring a Class Usage Profile	320
Configuring the File Information	321
Configuring the Interval	321
Creating a Class Usage Profile to Collect Source Class Usage Statistics ..	321
Creating a Class Usage Profile to Collect Destination Class Usage Statistics	322
Configuring the MIB Profile	322
Configuring the File Information	323
Configuring the Interval	323
Configuring the MIB Operation	323
Configuring MIB Object Names	324
Example: Configuring a MIB Profile	324
Configuring the Routing Engine Profile	325
Configuring Fields	325
Configuring the File Information	325
Configuring the Interval	326
Example: Configuring a Routing Engine Profile	326

Part 6

Chapter 19

Configuring Monitoring Options

Configuring Interface Alarms	329
Alarm Overview	329
Alarm Types	329
Alarm Severity	330

	Alarm Conditions	330
	Interface Alarm Conditions	331
	System Alarm Conditions	334
	Example: Configuring Interface Alarms	335
	Monitoring Active Alarms on a Device	338
	Monitoring Alarms	339
Chapter 20	Using RPM to Measure Network Performance	341
	RPM Overview	341
	RPM Probes	342
	RPM Tests	343
	Probe and Test Intervals	343
	Jitter Measurement with Hardware Timestamping	343
	RPM Statistics	344
	RPM Thresholds and Traps	345
	RPM for BGP Monitoring	345
	IPv6 RPM Probes	345
	Two-Way Active Measurement Protocol (TWAMP) Overview	346
	Benefits of TWAMP	347
	Guidelines for Configuring RPM Probes for IPv6	348
	RPM Support for VPN Routing and Forwarding	349
	Example: Configuring Basic RPM Probes	350
	Example: Configuring RPM Using TCP and UDP Probes	354
	Example: Configuring RPM Probes for BGP Monitoring	357
	Directing RPM Probes to Select BGP Devices	360
	Configuring IPv6 RPM Probes	360
	Tuning RPM Probes	361
	RPM Configuration Options	362
	Monitoring RPM Probes	366
Chapter 21	Configuring IP Monitoring	371
	IP Monitoring Overview	371
	Understanding IP Monitoring Test Parameters	372
	Example: Configuring IP Monitoring on SRX Series Devices	373
	Understanding IP Monitoring Through Redundant Ethernet Interface Link Aggregation Groups	376
	Example: Configuring IP Monitoring on SRX Series Devices	376
	Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring	382
Part 7	Monitoring Common Security Features	
Chapter 22	Displaying Real-Time Information from Device to Host	389
	Displaying Real-Time Monitoring Information	389
	Displaying Multicast Path Information	391

Chapter 23	Monitoring Application Layer Gateways Features	395
	Monitoring H.323 ALG Information	395
	Monitoring MGCP ALGs	396
	Monitoring MGCP ALG Calls	397
	Monitoring MGCP ALG Counters	397
	Monitoring MGCP ALG Endpoints	399
	Monitoring SCCP ALGs	399
	Monitoring SCCP ALG Calls	400
	Monitoring SCCP ALG Counters	400
	Monitoring SIP ALGs	402
	Monitoring SIP ALG Calls	402
	Monitoring SIP ALG Counters	402
	Monitoring SIP ALG Rate Information	404
	Monitoring SIP ALG Transactions	405
	Monitoring Voice ALG H.323	406
	Monitoring Voice ALG MGCP	408
	Monitoring Voice ALG SCCP	411
	Monitoring Voice ALG SIP	414
	Monitoring Voice ALG Summary	419
Chapter 24	Monitoring Interfaces and Switching Functions	421
	Displaying Real-Time Interface Information	421
	Monitoring Address Pools	423
	Monitoring Ethernet Switching	424
	Monitoring GVRP	425
	Monitoring Interfaces	426
	Monitoring MPLS Traffic Engineering Information	427
	Monitoring MPLS Interfaces	428
	Monitoring MPLS LSP Information	428
	Monitoring MPLS LSP Statistics	430
	Monitoring RSVP Session Information	431
	Monitoring MPLS RSVP Interfaces Information	432
	Monitoring PPP	433
	Monitoring PPPoE	434
	Monitoring Spanning Tree	438
	Monitoring the WAN Acceleration Interface	439
Chapter 25	Monitoring NAT	441
	Monitoring NAT	441
	Monitoring Source NAT Information	441
	Monitoring Destination NAT Information	447
	Monitoring Static NAT Information	449
	Monitoring Incoming Table Information	451
	Monitoring Interface NAT Port Information	452
Chapter 26	Monitoring Security Policies	453
	Monitoring Policy Statistics	453
	Monitoring Routing Information	454
	Monitoring Route Information	454
	Monitoring RIP Routing Information	456

	Monitoring OSPF Routing Information	458
	Monitoring BGP Routing Information	460
	Monitoring Security Events by Policy	462
	Monitoring Security Features	464
	Monitoring Policies	464
	Checking Policies	466
	Monitoring Screen Counters	469
	Monitoring IDP Status	472
	Monitoring Flow Gate Information	473
	Monitoring Firewall Authentication Table	474
	Monitoring Firewall Authentication History	476
	Monitoring 802.1x	477
Chapter 27	Monitoring Events, Services and System	479
	Monitoring DHCP Client Bindings	479
	Monitoring Events	480
	Monitoring the System	482
	Monitoring System Properties for SRX Series Devices	482
	Monitoring Chassis Information	485
	System Health Management for SRX Series Devices	486
Chapter 28	Monitoring Unified Threat Management Features	489
	Monitoring Antivirus Scan Engine Status	489
	Monitoring Antivirus Scan Results	490
	Monitoring Antivirus Session Status	492
	Monitoring Content Filtering Configurations	493
	Monitoring Reports	494
	Threats Monitoring Report	494
	Traffic Monitoring Report	499
	Monitoring Web Filtering Configurations	501
Chapter 29	Monitoring VPNs	503
	Monitoring VPNs	503
	Monitoring IKE Gateway Information	503
	Monitoring IPsec VPN—Phase I	507
	Monitoring IPsec VPN—Phase II	508
	Monitoring IPsec VPN Information	509

Part 8	Resource Monitoring of Memory Regions and Types Using CLI and SNMP Queries	
Chapter 30	Effective Troubleshooting of System Performance With Resource Monitoring Methodology	517
	Resource Monitoring Usage Computation Overview	517
	Resource Monitoring and Usage Computation For Trio-Based Line Cards	518
	Resource Monitoring and Usage Computation For I-Chip-Based Line Cards	518
	Diagnosing and Debugging System Performance by Configuring Memory	
	Resource Usage Monitoring on MX Series Routers	520
	Troubleshooting the Mismatch of jnxNatObjects Values for MS-DPC and MS-MIC	523
	Managed Objects for Ukernel Memory for a Packet Forwarding Engine in an FPC Slot	524
	Managed Objects for Packet Forwarding Engine Memory Statistics Data	524
	Managed Objects for Next-Hop, Jtree, and Firewall Filter Memory for a Packet Forwarding Engine in an FPC Slot	525
	jnxPfeMemoryErrorsTable	525
	pfeMemoryErrors	526
Part 9	Troubleshooting	
Chapter 31	Configuring Data Path Debugging and Trace Options	529
	Understanding Data Path Debugging for SRX Series Devices	529
	Debugging the Data Path (CLI Procedure)	530
	Example: Configuring End-to-End Debugging on SRX Series Device	531
	Understanding Security Debugging Using Trace Options	536
	Setting Security Trace Options (CLI Procedure)	537
	Displaying Log and Trace Files	538
	Displaying Output for Security Trace Options	538
	Displaying Multicast Trace Operations	539
	Using the J-Web Traceroute Tool	540
	J-Web Traceroute Results and Output Summary	542
	Understanding Flow Debugging Using Trace Options	543
	Setting Flow Debugging Trace Options (CLI Procedure)	543
	Displaying a List of Devices	544
Chapter 32	Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits	547
	MPLS Connection Checking Overview	547
	Understanding Ping MPLS	549
	MPLS Enabled	549
	Loopback Address	549
	Source Address for Probes	550
	Using the ping Command	550
	Using the J-Web Ping Host Tool	552
	J-Web Ping Host Results and Output Summary	554
	Using the J-Web Ping MPLS Tool	555
	J-Web Ping MPLS Results and Output Summary	558

	Pinging Layer 2 Circuits	559
	Pinging Layer 2 VPNs	560
	Pinging Layer 3 VPNs	562
	Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs	563
Chapter 33	Using Packet Capture to Analyze Network Traffic	565
	Packet Capture Overview	565
	Packet Capture on Device Interfaces	566
	Firewall Filters for Packet Capture	567
	Packet Capture Files	567
	Analysis of Packet Capture Files	567
	Example: Enabling Packet Capture on a Device	568
	Example: Configuring Packet Capture on an Interface	572
	Example: Configuring a Firewall Filter for Packet Capture	573
	Example: Configuring Packet Capture for Datapath Debugging	575
	Disabling Packet Capture	579
	Deleting Packet Capture Files	579
	Changing Encapsulation on Interfaces with Packet Capture Configured	580
	Displaying Packet Headers	581
	Using the J-Web Packet Capture Tool	586
	J-Web Packet Capture Results and Output Summary	589
Chapter 34	Troubleshooting Security Devices	591
	Recovering the Root Password for SRX Series Devices	591
	Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only)	593
	Troubleshooting the Link Services Interface	593
	Determine Which CoS Components Are Applied to the Constituent Links	594
	Determine What Causes Jitter and Latency on the Multilink Bundle	596
	Determine If LFI and Load Balancing Are Working Correctly	596
	Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device	603
	Troubleshooting Security Policies	604
	Synchronizing Policies Between Routing Engine and Packet Forwarding Engine	604
	Checking a Security Policy Commit Failure	604
	Verifying a Security Policy Commit	605
	Debugging Policy Lookup	606
	Understanding Log Error Messages for Troubleshooting ISSU-Related Problems	606
	Chassisd Process Errors	606
	Kernel State Synchronization	607
	Installation-Related Errors	607
	ISSU Support-Related Errors	607
	Redundancy Group Failover Errors	608
	Initial Validation Checks Fail	608
	Understanding Common Error Handling for ISSU	609

Part 10

Chapter 35

Configuration Statements and Operational Commands

Configuration Statements: Accounting Options, Source Class Usage and Destination Class Usage Options 615

accounting-options	617
allow-clear (Accounting Options)	619
archive-sites	620
backup-on-failure (Accounting Options)	621
class-usage-profile	622
cleanup-interval (Accounting Options)	623
compress (Accounting Options)	624
counters	624
destination-classes	625
egress-stats (Flat-File Accounting Options)	626
fields (Flat-File Accounting Options)	628
fields (for Interface Profiles)	630
fields (for Routing Engine Profiles)	631
file (Associating with a Profile)	632
file (Configuring a Log File)	633
file (Flat-File Accounting Options)	634
files	635
filter-profile	636
flat-file-profile (Accounting Options)	637
format (Flat-File Accounting Options)	639
general-param (Flat-File Accounting Options)	640
ingress-stats (Flat-File Accounting Options)	642
interface-profile	643
interval	644
interval (Flat-File Accounting Options)	645
l2-stats (Flat-File Accounting Options)	646
mib-profile	647
mpls (Security Forwarding Options)	648
nonpersistent	649
object-names	650
operation	651
overall-packet (Flat-File Accounting Options)	652
push-backup-to-master (Accounting Options)	653
routing-engine-profile	654
schema-version (Flat-File Accounting Options)	655
size	656
source-classes	657
start-time	658
traceoptions (System Accounting)	659
transfer-interval	661

Chapter 36

Configuration Statements: Chassis Cluster 663

cluster (Chassis)	664
global-threshold	665
global-weight	666
ip-monitoring	667

	ip-monitoring (Services)	669
	next-hop	670
Chapter 37	Configuration Statements: Datapath Debug	671
	action-profile	672
	capture-file (Security)	674
	datapath-debug	675
	flow (Security Flow)	677
	icmp	679
	maximum-capture-size (Datapath Debug)	679
	traceoptions (Security Datapath Debug)	680
Chapter 38	Configuration Statements: Health Monitoring	683
	falling-threshold	684
	health-monitor	685
	idp (SNMP)	686
	routing-engine (SNMP Resource Level)	687
	interval	688
	rising-threshold	689
Chapter 39	Configuration Statements: Remote Monitoring (RMON)	691
	alarm (SNMP RMON)	692
	community	693
	description	694
	event	695
	falling-event-index	696
	falling-threshold	697
	falling-threshold-interval	698
	interval	699
	request-type	700
	rising-event-index	701
	rising-threshold	702
	rmon	702
	sample-type	703
	startup-alarm	704
	syslog-subtag	705
	type	706
	variable	707
Chapter 40	Configuration Statements: Resource Monitoring for Memory Regions	709
	free-fw-memory-watermark (Resource Monitor)	710
	free-heap-memory-watermark (Resource Monitor)	711
	free-nh-memory-watermark (Resource Monitor)	712
	high-cos-queue-threshold	713
	high-threshold (Resource Monitor)	714
	no-load-throttle (Resource Monitor)	714
	no-logging (Resource Monitor)	715
	no-throttle (Resource Monitor)	716
	resource-category jtree (Resource Monitor)	717
	resource-monitor	718

	subscribers-limit (Resource Monitor)	720
	traceoptions (Resource Monitor)	721
Chapter 41	Configuration Statements: Security Alarms	723
	decryption-failures	723
	idp (Security Alarms)	724
Chapter 42	Configuration Statements: SNMP	725
	access-list	726
	agent-address	727
	alarm-id	728
	alarm-list-name	729
	alarm-management	730
	alarm-state	731
	authorization	732
	categories	733
	client-list	734
	client-list-name	735
	clients	736
	commit-delay	737
	community (SNMP)	738
	contact (SNMP)	739
	description	740
	destination-port	741
	enterprise-oid	741
	filter-duplicates	742
	filter-interfaces	743
	interface (SNMP)	744
	location (SNMP)	745
	logical-system	746
	logical-system-trap-filter	747
	name	747
	nonvolatile	748
	oid	749
	proxy (snmp)	750
	routing-instance	752
	routing-instance-access	753
	snmp	753
	snmp-value-match-msmic (Services NAT Options)	754
	source-address	755
	targets	756
	traceoptions (SNMP)	757
	trap-group	759
	trap-options	760
	version (SNMP)	761
	view (Associating a MIB View with a Community)	762
	view (Configuring a MIB View)	763

Chapter 43	Configuration Statements: SNMPv3	765
	address	767
	address-mask	767
	authentication-key	768
	authentication-md5	769
	authentication-none	770
	authentication-password	771
	authentication-sha	772
	community-name	773
	context (SNMPv3)	774
	engine-id	775
	group (Configuring Group Name)	776
	group (Defining Access Privileges for an SNMPv3 Group)	777
	retry-count	778
	timeout	779
	local-engine	780
	message-processing-model	781
	notify	782
	notify-filter (Applying to the Management Target)	783
	notify-filter (Configuring the Profile Name)	784
	notify-view	785
	oid	786
	parameters	787
	port	788
	privacy-3des	789
	privacy-aes128	790
	privacy-des	791
	privacy-key	792
	privacy-none	793
	privacy-password	794
	read-view	795
	remote-engine	796
	routing-instance	797
	security-level (Defining Access Privileges)	798
	security-level (Generating SNMP Notifications)	799
	security-model (Access Privileges)	800
	security-model (Group)	801
	security-model (SNMP Notifications)	802
	security-name (Community String)	803
	security-name (Security Group)	804
	security-name (SNMP Notifications)	805
	security-to-group	806
	snmp-community	807
	tag	808
	tag-list	809
	target-address	810
	target-parameters	811
	type	812
	user	813

	usm	814
	v3	816
	vacm	819
	write-view	820
Chapter 44	Configuration Statements: TWAMP	821
	twamp-server	822
	twamp-client	824
Chapter 45	Operational Commands	827
	clear chassis cluster ip-monitoring failure-count	829
	clear chassis cluster ip-monitoring failure-count ip-address	830
	clear ilmi statistics	831
	clear interfaces statistics	832
	clear services rpm twamp server connection	833
	clear snmp history	834
	clear snmp statistics	835
	request pppoe connect	837
	request pppoe disconnect	838
	request services ip-monitoring preempt-restore policy	839
	request services rpm twamp start	840
	request services rpm twamp stop	841
	request snmp spoof-trap	842
	show chassis alarms	848
	show chassis cluster ip-monitoring status redundancy-group	850
	show interfaces (SRX Series)	853
	show interfaces snmp-index	885
	show interfaces summary	886
	show ilmi statistics	888
	show security alarms	891
	show security datapath-debug capture	895
	show security datapath-debug counter	896
	show security monitoring	897
	show security monitoring fpc fpc-number	899
	show security monitoring performance session	902
	show security monitoring performance spu	903
	show services ip-monitoring status	905
	show services rpm twamp client connection	909
	show services rpm twamp client history-results	911
	show services rpm twamp client probe-results	913
	show services rpm twamp client session	918
	show services rpm twamp server connection	920
	show services rpm twamp server session	922
	show snmp health-monitor	924
	show snmp inform-statistics	931
	show snmp mib	933
	show snmp rmon	936
	show snmp statistics	940
	show snmp stats-response-statistics	948
	show snmp v3	950

show system alarms	953
show system alarms	954
show system resource-monitor fpc	957

List of Figures

Part 2	Network Monitoring Using SNMP	
Chapter 7	Configuring SNMPv3	115
	Figure 1: Inform Request and Response	144
Chapter 8	Configuring SNMP for Routing Instances	155
	Figure 2: SNMP Data for Routing Instances	156
Part 3	Remote Monitoring (RMON) with SNMP	
Chapter 15	Using RMON to Monitor Network Service Quality	257
	Figure 3: Setting Thresholds	258
	Figure 4: Network Entry Points	261
	Figure 5: Regional Points of Presence	263
	Figure 6: Measurements to Each Router	263
	Figure 7: Network Behavior During Congestion	278
Part 6	Configuring Monitoring Options	
Chapter 20	Using RPM to Measure Network Performance	341
	Figure 8: TWAMP Test	347
	Figure 9: Four Elements of TWAMP	347
	Figure 10: Sample RPM Graphs	367
Chapter 21	Configuring IP Monitoring	371
	Figure 11: IP Monitoring on an SRX Series Device Topology Example	377
Part 9	Troubleshooting	
Chapter 34	Troubleshooting Security Devices	591
	Figure 12: PPP and MLPPP Headers	599

List of Tables

	About the Documentation	xxxi
	Table 1: Notice Icons	xxxiv
	Table 2: Text and Syntax Conventions	xxxiv
Part 1	Overview	
Chapter 1	Network Management Overview	3
	Table 3: Device Management Features in Junos OS	4
Chapter 2	Introduction to Network Monitoring	7
	Table 4: J-Web Interface Troubleshoot Options	9
	Table 5: CLI Diagnostic Command Summary	10
Part 2	Network Monitoring Using SNMP	
Chapter 4	SNMP MIBs and Traps Supported by Junos OS	21
	Table 6: Enterprise-specific MIBs supported by Junos OS	21
	Table 7: Standard MIBs supported by Junos OS	32
	Table 8: Standard Supported SNMP Version 1 Traps	50
	Table 9: Standard Supported SNMP Version 2 Traps	53
	Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps	58
	Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps	64
Chapter 7	Configuring SNMPv3	115
	Table 12: Values to Use in Example	146
Chapter 8	Configuring SNMP for Routing Instances	155
	Table 13: MIB Support for Routing Instances (Juniper Networks MIBs)	156
	Table 14: Class 1 MIB Objects (Standard and Juniper MIBs)	160
	Table 15: Class 2 MIB Objects (Standard and Juniper MIBs)	164
	Table 16: Class 3 MIB Objects (Standard and Juniper MIBs)	165
	Table 17: Class 4 MIB Objects (Standard and Juniper MIBs)	166
Chapter 9	Configuring SNMP Remote Operations	173
	Table 18: Results in pingProbeHistoryTable: After the First Ping Test	181
	Table 19: Results in pingProbeHistoryTable: After the First Probe of the Second Test	182
	Table 20: Results in pingProbeHistoryTable: After the Second Ping Test	182
	Table 21: traceRouteProbeHistoryTable	191
Chapter 10	Tracing SNMP Activity	195

	Table 22: SNMP Tracing Flags	204
Chapter 11	SNMP FAQs	207
	Table 23: Monitored Object Instances	215
Part 3	Remote Monitoring (RMON) with SNMP	
Chapter 15	Using RMON to Monitor Network Service Quality	257
	Table 24: RMON Event Table	259
	Table 25: RMON Alarm Table	259
	Table 26: jnxRmon Alarm Extensions	260
	Table 27: Real-Time Performance Monitoring Configuration Options	266
	Table 28: Health Metrics	268
	Table 29: Counter Values for vlan-ccc Encapsulation	274
	Table 30: Performance Metrics	275
	Table 31: Inbound Traffic Per Class	278
	Table 32: Inbound Counters	279
	Table 33: Outbound Counters for ATM Interfaces	279
	Table 34: Outbound Counters for Non-ATM Interfaces	280
	Table 35: Dropped Traffic Counters	280
Part 4	Health Monitoring with SNMP	
Chapter 16	Configuring Health Monitoring	285
	Table 36: Monitored Object Instances	286
Part 5	Gathering Statistics for Accounting Purposes Using Accounting Options, Source Class Usage and Destination Class Usage Options	
Chapter 17	Accounting Options, Source Class Usage and Destination Class Usage Options Overview	291
	Table 37: Types of Accounting Profiles	291
Part 6	Configuring Monitoring Options	
Chapter 19	Configuring Interface Alarms	329
	Table 38: Interface Alarm Conditions	331
	Table 39: System Alarm Conditions and Corrective Actions	334
	Table 40: Alarms Monitoring Page	339
Chapter 20	Using RPM to Measure Network Performance	341
	Table 41: RPM Statistics	344
	Table 42: RPM Configuration Summary	362
	Table 43: Summary of Key RPM Output Fields	367
Chapter 21	Configuring IP Monitoring	371
	Table 44: Test Parameters and Default Values	372
	Table 45: Threshold Supported and Description	373

Part 7	Monitoring Common Security Features	
Chapter 22	Displaying Real-Time Information from Device to Host	389
	Table 46: CLI traceroute monitor Command Options	389
	Table 47: CLI traceroute monitor Command Output Summary	390
	Table 48: CLI mtrace from-source Command Options	391
	Table 49: CLI mtrace from-source Command Output Summary	393
Chapter 23	Monitoring Application Layer Gateways Features	395
	Table 50: Summary of Key H.323 Counters Output Fields	395
	Table 51: Summary of Key MGCP Calls Output Fields	397
	Table 52: Summary of Key MGCP Counters Output Fields	398
	Table 53: Summary of Key MGCP Endpoints Output Fields	399
	Table 54: Summary of Key SCCP Calls Output Fields	400
	Table 55: Summary of Key SCCP Counters Output Fields	400
	Table 56: Summary of Key SIP Calls Output Fields	402
	Table 57: Summary of Key SIP Counters Output Fields	403
	Table 58: Summary of Key SIP Rate Output Fields	405
	Table 59: Summary of Key SIP Transactions Output Fields	405
	Table 60: ALG H.323 Monitoring Page	406
	Table 61: Voice ALG MGCP Monitoring Page	409
	Table 62: Voice ALG SCCP Monitoring Page	411
	Table 63: Voice ALG SIP Monitoring Page	414
	Table 64: Voice ALG Summary Monitoring Page	419
Chapter 24	Monitoring Interfaces and Switching Functions	421
	Table 65: CLI monitor interface Output Control Keys	422
	Table 66: CLI monitor interface traffic Output Control Keys	422
	Table 67: Address Pools Monitoring Page	423
	Table 68: Summary of Ethernet Switching Output Fields	425
	Table 69: GVRP Monitoring Page	426
	Table 70: Summary of Key MPLS Interface Information Output Fields	428
	Table 71: Summary of Key MPLS LSP Information Output Fields	429
	Table 72: Summary of Key MPLS LSP Statistics Output Fields	430
	Table 73: Summary of Key RSVP Session Information Output Fields	431
	Table 74: Summary of Key RSVP Interfaces Information Output Fields	432
	Table 75: Summary of Key PPPoE Output Fields	434
	Table 76: Spanning Tree Monitoring Page	438
Chapter 25	Monitoring NAT	441
	Table 77: Source NAT Monitoring Page	441
	Table 78: Summary of Key Destination NAT Output Fields	447
	Table 79: Summary of Key Static NAT Output Fields	449
	Table 80: Summary of Key Incoming Table Output Fields	451
	Table 81: Summary of Key Interface NAT Output Fields	452
Chapter 26	Monitoring Security Policies	453
	Table 82: Filtering Route Messages	455
	Table 83: Summary of Key Routing Information Output Fields	455
	Table 84: Summary of Key RIP Routing Output Fields	457
	Table 85: Summary of Key OSPF Routing Output Fields	458

	Table 86: Summary of Key BGP Routing Output Fields	460
	Table 87: View Policy Log Fields	462
	Table 88: Policy Events Detail Fields	463
	Table 89: Security Policies Monitoring Output Fields	464
	Table 90: Check Policies Output	467
	Table 91: Summary of Key Screen Counters Output Fields	469
	Table 92: Summary of IDP Status Output Fields	472
	Table 93: Summary of Key Flow Gate Output Fields	473
	Table 94: Summary of Key Firewall Authentication Table Output Fields	474
	Table 95: Summary of Key Firewall Authentication History Output Fields	476
	Table 96: Summary of Dot1X Output Fields	478
Chapter 27	Monitoring Events, Services and System	479
	Table 97: Summary of Key DHCP Client Binding Output Fields	479
	Table 98: Events Monitoring Page	480
Chapter 28	Monitoring Unified Threat Management Features	489
	Table 99: Statistics Tab Output in the Threats Report	494
	Table 100: Activities Tab Output in the Threats Report	497
	Table 101: Traffic Report Output	499
Chapter 29	Monitoring VPNs	503
	Table 102: Summary of Key IKE SA Information Output Fields	503
	Table 103: IPsec VPN—Phase I Monitoring Page	507
	Table 104: IPsec VPN—Phase II Monitoring Page	508
	Table 105: Summary of Key IPsec VPN Information Output Fields	510
Part 8	Resource Monitoring of Memory Regions and Types Using CLI and SNMP Queries	
Chapter 30	Effective Troubleshooting of System Performance With Resource Monitoring Methodology	517
	Table 106: jnxPfeMemoryUKernTable	524
	Table 107: jnxPfeMemory Table	524
	Table 108: jnxPfeMemoryForwardingTable	525
	Table 109: jnxPfeMemoryErrorsTable	526
	Table 110: pfeMemoryErrors	526
Part 9	Troubleshooting	
Chapter 31	Configuring Data Path Debugging and Trace Options	529
	Table 111: CLI mtrace monitor Command Output Summary	539
	Table 112: Traceroute Field Summary	541
	Table 113: J-Web Traceroute Results and Output Summary	542
	Table 114: CLI traceroute Command Options	544
Chapter 32	Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits	547
	Table 115: Options for Checking MPLS Connections	548
	Table 116: CLI ping Command Options	550
	Table 117: J-Web Ping Host Field Summary	553
	Table 118: Ping Host Results and Output	554

	Table 119: J-Web Ping MPLS Field Summary	556
	Table 120: J-Web Ping MPLS Results and Output Summary	558
	Table 121: CLI ping mpls l2circuit Command Options	560
	Table 122: CLI ping mpls l2vpn Command Options	561
	Table 123: CLI ping mpls l3vpn Command Options	562
	Table 124: CLI ping mpls ldp and ping mpls lsp-end-point Command Options	563
Chapter 33	Using Packet Capture to Analyze Network Traffic	565
	Table 125: CLI monitor traffic Command Options	582
	Table 126: CLI monitor traffic Match Conditions	583
	Table 127: CLI monitor traffic Logical Operators	585
	Table 128: CLI monitor traffic Arithmetic, Binary, and Relational Operators	585
	Table 129: Packet Capture Field Summary	587
	Table 130: J-Web Packet Capture Results and Output Summary	590
Chapter 34	Troubleshooting Security Devices	591
	Table 131: CoS Components Applied on Multilink Bundles and Constituent Links	594
	Table 132: PPP and MLPPP Encapsulation Overhead	599
	Table 133: Number of Packets Transmitted on a Queue	602
	Table 134: ISSU-Related Errors and Solutions	609
Part 10	Configuration Statements and Operational Commands	
Chapter 45	Operational Commands	827
	Table 135: show chassis alarms Output Fields	848
	Table 136: show chassis cluster ip-monitoring status Output Fields	850
	Table 137: show chassis cluster ip-monitoring status redundancy group Reason Fields	851
	Table 138: show interfaces Output Fields	856
	Table 139: show interfaces summary Output Fields	886
	Table 140: show ilmi statistics Output Fields	889
	Table 141: show security alarms	892
	Table 142: show security monitoring fpc fpc-number Output Fields	899
	Table 143: show services ip-monitoring status Output Fields	905
	Table 144: show services rpm twamp client connection Output Fields	909
	Table 145: show services rpm twamp client history-results Output Fields	911
	Table 146: show services twamp client probe-results Output Fields	913
	Table 147: show services rpm twamp client session Output Fields	918
	Table 148: show services rpm twamp server connection Output Fields	920
	Table 149: show services rpm twamp server session Output Fields	922
	Table 150: show snmp health-monitor Output Fields	924
	Table 151: show snmp inform-statistics Output Fields	931
	Table 152: show snmp mib Output Fields	934
	Table 153: show snmp rmon Output Fields	936
	Table 154: show snmp statistics Output Fields	941
	Table 155: show snmp statistics subagents Output Fields	944
	Table 156: show snmp stats-response-statistics Output Fields	948
	Table 157: show snmp v3 Output Fields	951
	Table 158: show system alarms Output Fields	955

Table 159: show system resource-monitor fpc Output Fields 958

About the Documentation

- [Documentation and Release Notes on page xxxi](#)
- [Supported Platforms on page xxxi](#)
- [Using the Examples in This Manual on page xxxii](#)
- [Documentation Conventions on page xxxiii](#)
- [Documentation Feedback on page xxxv](#)
- [Requesting Technical Support on page xxxvi](#)

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- [ACX Series](#)
- [M Series](#)
- [MX Series](#)
- [T Series](#)
- [PTX Series](#)
- [SRX Series](#)
- [vSRX](#)
- [QFX Series](#)
- [EX Series](#)

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {  
    file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

```
[edit]  
user@host# edit system scripts  
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]  
user@host# load merge relative /var/tmp/ex-script-snippet.conf  
load complete
```

For more information about the **load** command, see [CLI Explorer](#).

Documentation Conventions

[Table 1 on page xxxiv](#) 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 xxxiv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options { static { route default { nexthop <i>address</i>; retain; } } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

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 post-sales 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 <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://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: <http://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <http://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- 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 <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

- [Network Management Overview on page 3](#)
- [Introduction to Network Monitoring on page 7](#)

CHAPTER 1

Network Management Overview

- [Understanding Device Management Functions in Junos OS on page 3](#)
- [Understanding the Integrated Local Management Interface on page 6](#)

Understanding Device Management Functions in Junos OS

Supported Platforms [ACX Series, M Series, MX Series, T Series](#)

After you have installed a device into your network, you need to manage the device within your network. Device management can be divided into five tasks:

- Fault management—Monitor the device; detect and fix faults.
- Configuration management—Configure device attributes.
- Accounting management—Collect statistics for accounting purposes.
- Performance management—Monitor and adjust device performance.
- Security management—Control device access and authenticate users.

The Junos[®] operating system (Junos OS) network management features work in conjunction with an operations support system (OSS) to manage the devices within the network. Junos OS can assist you in performing these management tasks, as described in [Table 3 on page 4](#).

Table 3: Device Management Features in Junos OS

Task	Junos OS Feature
Fault management	<p>Monitor and see faults using:</p> <ul style="list-style-type: none">• Operational mode commands—For more information about operational mode commands, see the CLI Explorer.• SNMP MIBs—For more information about SNMP MIBs supported by Junos OS, see ““Standard SNMP MIBs Supported by Junos OS” on page 32” and ““Enterprise-Specific SNMP MIBs Supported by Junos OS” on page 21.• Standard SNMP traps—For more information about standard SNMP traps, see the “Standard SNMP Traps Supported by Junos OS” on page 50.• Enterprise-specific SNMP traps—For more information about enterprise-specific traps, see ““Enterprise-Specific SNMP Traps Supported by Junos OS” on page 57”.• System log messages—For more information about how to configure system log messages, see the <i>Junos OS Administration Library</i>. For more information about how to view system log messages, see the System Log Explorer.
Configuration management	<ul style="list-style-type: none">• Configure router attributes using the command-line interface (CLI), the Junos XML management protocol, and the NETCONF XML management protocol. For more information about configuring the router using the CLI, see the <i>Junos OS Administration Library</i>. For more information about configuring the router using the APIs, see the <i>Junos XML Management Protocol Guide</i> and <i>NETCONF XML Management Protocol Guide</i>.• Configuration Management MIB—For more information about the Configuration Management MIB, see Configuration Management MIB.

Table 3: Device Management Features in Junos OS (*continued*)

Task	Junos OS Feature
Accounting management	<p>Perform the following accounting-related tasks:</p> <ul style="list-style-type: none"> Collect statistics for interfaces, firewall filters, destination classes, source classes, and the Routing Engine. For more information about collecting statistics, see “Accounting Options Configuration” on page 296. Use interface-specific traffic statistics and other counters, available in the Standard Interfaces MIB, Juniper Networks enterprise-specific extensions to the Interfaces MIB, and media-specific MIBs, such as the enterprise-specific ATM MIB. Use per-ATM virtual circuit (VC) counters, available in the enterprise-specific ATM MIB. For more information about the ATM MIB, see ATM MIB. Group source and destination prefixes into source classes and destination classes and count packets for those classes. Collect destination class and source class usage statistics. For more information about classes, see “Destination Class Usage MIB” and “Source Class Usage MIB”, “Configuring Class Usage Profiles” on page 320, the <i>Junos OS Network Interfaces Library for Routing Devices</i>, and the <i>Junos OS Routing Protocols Library</i>. Count packets as part of a firewall filter. For more information about firewall filter policies, see “Enterprise-Specific SNMP MIBs Supported by Junos OS” on page 21 and the <i>Junos OS Routing Protocols Library</i>. Sample traffic, collect the samples, and send the collection to a host running the CAIDA cflowd utility. For more information about CAIDA and cflowd, see the <i>Junos OS Routing Protocols Library for Security Devices</i>.
Performance management	<p>Monitor performance in the following ways:</p> <ul style="list-style-type: none"> Use operational mode commands. For more information about monitoring performance using operational mode commands, see the CLI Explorer. Use firewall filter. For more information about performance monitoring using firewall filters, see the <i>Junos OS Routing Protocols Library</i>. Sample traffic, collect the samples, and send the samples to a host running the CAIDA cflowd utility. For more information about CAIDA and cflowd, see the <i>Junos OS Routing Protocols Library</i>. Use the enterprise-specific Class-of-Service MIB. For more information about this MIB, see Class-of-Service MIB.
Security management	<p>Assure security in your network in the following ways:</p> <ul style="list-style-type: none"> Control access to the router and authenticate users. For more information about access control and user authentication, see the <i>Junos OS Administration Library</i>. Control access to the router using SNMPv3 and SNMP over IPv6. For more information, see “Configuring the Local Engine ID” on page 120 and “Tracing SNMP Activity on a Device Running Junos OS” on page 201.

- Related Documentation**
- [Understanding the Integrated Local Management Interface on page 6](#)
 - [Understanding the SNMP Implementation in Junos OS](#)
 - [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 261](#)
 - [Accounting Options Overview on page 291](#)

Understanding the Integrated Local Management Interface

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

The Integrated Local Management Interface (ILMI) provides a mechanism for Asynchronous Transfer Mode (ATM)-attached devices, such as hosts, routers, and ATM switches, to transfer management information. ILMI provides bidirectional exchange of management information between two ATM interfaces across a physical connection. ILMI information is exchanged over a direct encapsulation of SNMP version 1 (RFC 1157, *A Simple Network Management Protocol*) over ATM Adaptation Layer 5 (AAL5) using a virtual path identifier/virtual channel identifier (VPI/VCI) value (VPI=0, VCI=16).

Junos OS supports only two ILMI MIB variables: **atmfMYIPNmAddress** and **atmfPortMyIfname**. For ATM1 and ATM2 intelligent queuing (IQ) interfaces, you can configure ILMI to communicate directly with an attached ATM switch to enable querying of the switch's IP address and port number.

For more information about the ILMI MIB, see **atmfMYIPNmAddress** or **atmfPortMyIfname** in the [SNMP MIB Explorer](#).

- Related Documentation**
- [Understanding Device Management Functions in Junos OS on page 3](#)

CHAPTER 2

Introduction to Network Monitoring

- [Monitoring Overview on page 7](#)
- [Diagnostic Tools Overview on page 8](#)

Monitoring Overview

Supported Platforms [SRX Series, vSRX](#)

Junos OS supports a suite of J-Web tools and CLI operational mode commands for monitoring the system health and performance of your device. Monitoring tools and commands display the current state of the device. To use the J-Web user interface and CLI operational tools, you must have the appropriate access privileges.

You can use the J-Web Monitor option to monitor a device. J-Web results appear in the browser.

You can also monitor the device with CLI operational mode commands. CLI command output appears on the screen of your console or management device, or you can filter the output to a file. For operational commands that display output, such as the **show** commands, you can redirect the output into a filter or a file. When you display help about these commands, one of the options listed is **|**, called a *pipe*, which allows you to filter the command output.

For example, if you enter the **show configuration** command, the complete device configuration appears on the screen. To limit the display to only those lines of the configuration that contain **address**, enter the **show configuration** command using a pipe into the **match** filter:

```
user@host> show configuration | match address
address-range low 192.168.3.2 high 192.168.3.254;
address-range low 192.168.71.71 high 192.168.71.254;
address 192.168.71.70/21;
address 192.168.2.1/24;
address 127.0.0.1/32;
```

For a complete list of the filters, type a command, followed by the pipe, followed by a question mark (?):

```
user@host> show configuration | ?
```

Possible completions:

compare	Compare configuration changes with prior version
count	Count occurrences
display	Show additional kinds of information
except	Show only text that does not match a pattern
find	Search for first occurrence of pattern
hold	Hold text without exiting the prompt
last	Display end of output only
match	Show only text that matches a pattern
no-more	Don't paginate output
request	Make system-level requests
resolve	Resolve IP addresses
save	Save output text to file
trim	Trim specified number of columns from start of line

You can specify complex expressions as an option for the **match** and **except** filters.



NOTE: To filter the output of configuration mode commands, use the filter commands provided for the operational mode commands. In configuration mode, an additional filter is supported.

- Related Documentation**
- [Monitoring Interfaces on page 426](#)
 - [Diagnostic Tools Overview on page 8](#)

Diagnostic Tools Overview

Supported Platforms [SRX Series, vSRX](#)

Juniper Networks devices support a suite of J-Web tools and CLI operational mode commands for evaluating system health and performance. Diagnostic tools and commands test the connectivity and reachability of hosts in the network.

- Use the J-Web Diagnose options to diagnose a device. J-Web results appear in the browser.
- Use CLI operational mode commands to diagnose a device. CLI command output appears on the screen of your console or management device, or you can filter the output to a file.

To use the J-Web user interface and CLI operational tools, you must have the appropriate access privileges.

This section contains the following topics:

- [J-Web Diagnostic Tools on page 9](#)
- [CLI Diagnostic Commands on page 9](#)

J-Web Diagnostic Tools

The J-Web diagnostic tools consist of the options that appear when you select **Troubleshoot** and **Maintain** in the task bar. [Table 4 on page 9](#) describes the functions of the Troubleshoot options.

Table 4: J-Web Interface Troubleshoot Options

Option	Function
Troubleshoot Options	
Ping Host	Allows you to ping a remote host. You can configure advanced options for the ping operation.
Ping MPLS	Allows you to ping an MPLS endpoint using various options.
Traceroute	Allows you to trace a route between the device and a remote host. You can configure advanced options for the traceroute operation.
Packet Capture	Allows you to capture and analyze router control traffic.
Maintain Options	
Files	Allows you to manage log, temporary, and core files on the device.
Upgrade	Allows you to upgrade and manage Junos OS packages.
Licenses	Displays a summary of the licenses needed and used for each feature that requires a license. Allows you to add licenses.
Reboot	Allows you to reboot the device at a specified time.

CLI Diagnostic Commands

The CLI commands available in operational mode allow you to perform the same monitoring, troubleshooting, and management tasks you can perform with the J-Web user interface. Instead of invoking the tools through a graphical interface, you use operational mode commands to perform the tasks.

You can perform certain tasks only through the CLI. For example, you can use the **mtrace** command to display trace information about a multicast path from a source to a receiver, which is a feature available only through the CLI.

To view a list of top-level operational mode commands, type a question mark (?) at the command-line prompt.

At the top level of operational mode are the broad groups of CLI diagnostic commands listed in [Table 5 on page 10](#).

Table 5: CLI Diagnostic Command Summary

Command	Function
Controlling the CLI Environment	
set option	Configures the CLI display.
Diagnosis and Troubleshooting	
clear	Clears statistics and protocol database information.
mtrace	Traces information about multicast paths from source to receiver.
monitor	Performs real-time debugging of various Junos OS components, including the routing protocols and interfaces.
ping	Determines the reachability of a remote network host.
ping mpls	Determines the reachability of an MPLS endpoint using various options.
test	Tests the configuration and application of policy filters and AS path regular expressions.
tracert	Traces the route to a remote network host.
Connecting to Other Network Systems	
ssh	Opens secure shell connections.
telnet	Opens Telnet sessions to other hosts on the network.
Management	
copy	Copies files from one location on the device to another, from the device to a remote system, or from a remote system to the device.
restart option	Restarts the various system processes, including the routing protocol, interface, and SNMP processes.
request	Performs system-level operations, including stopping and rebooting the device and loading Junos OS images.
start	Exits the CLI and starts a UNIX shell.
configuration	Enters configuration mode.
quit	Exits the CLI and returns to the UNIX shell.

- Related Documentation**
- [MPLS Connection Checking Overview on page 547](#)
 - [Understanding Ping MPLS on page 549](#)
 - [Using the J-Web Ping Host Tool on page 552](#)

- [Using the ping Command on page 550](#)

PART 2

Network Monitoring Using SNMP

- [SNMP Overview on page 15](#)
- [SNMP MIBs and Traps Supported by Junos OS on page 21](#)
- [Loading MIB Files to a Network Management System on page 73](#)
- [Configuring SNMP on page 77](#)
- [Configuring SNMPv3 on page 115](#)
- [Configuring SNMP for Routing Instances on page 155](#)
- [Configuring SNMP Remote Operations on page 173](#)
- [Tracing SNMP Activity on page 195](#)
- [SNMP FAQs on page 207](#)

CHAPTER 3

SNMP Overview

- [Understanding SNMP Implementation in Junos OS on page 15](#)
- [SNMPv3 Overview on page 18](#)

Understanding SNMP Implementation in Junos OS

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Do you use a central network management system (NMS)? Most NMS's use a version of Simple Network Management Protocol (SNMP) that can monitor the status of Junos OS devices that send unsolicited messages called traps. You can configure the IP address of your NMS so that Junos OS can send its traps.

SNMP uses a very basic form of authentication called community strings to control access between a manager and remote agents. Community strings are administrative names used to group collections of devices (and the agents running on them) into common management domains. If a manager and an agent share the same community, they can talk to one another.

Many people associate SNMP community strings with passwords and keys because the jobs they do are similar. As a result, SNMP communities are traditionally referred to as strings. The community string is the first level of management authentication implemented by the SNMP agent in Junos OS.

You might also want to configure remote logging on your device. Junos OS uses a system log (syslog) mechanism similar to many Unix devices to forward log messages to a specified log host address. This allows each of your devices to forward their messages to one central host, making it easier to monitor the network as a whole. Syslog is a very flexible and rich way of logging messages and is used by many device vendors to supplement the information provided by SNMP traps.

A typical SNMP implementation includes three components:

- Managed device
- SNMP agent
- Network management system (NMS)

A managed device is any device on a network, also known as a network element, that is managed by the network management system. Routers and switches are common examples of managed devices. The SNMP agent is the SNMP process that resides on the managed device and communicates with the network management system. The NMS is a combination of hardware and software that is used to monitor and administer a network.

The SNMP data is stored in a highly-structured, hierarchical format known as a management information base (MIB). The MIB structure is based on a tree structure, which defines a grouping of objects into related sets. Each object in the MIB is associated with an object identifier (OID), which names the object. The “leaf” in the tree structure is the actual managed object instance, which represents a resource, event, or activity that occurs in your network device.

The SNMP agent exchanges network management information with SNMP manager software running on an NMS, or host. The agent responds to requests for information and actions from the manager. The agent also controls access to the agent's MIB, the collection of objects that can be viewed or changed by the SNMP manager.

The SNMP manager collects information about network connectivity, activity, and events by polling managed devices.

Communication between the agent and the manager occurs in one of the following forms:

- **Get, GetBulk, and GetNext** requests—The manager requests information from the agent. The agent returns the information in a **Get** response message.
- **Set** requests—The manager changes the value of a MIB object controlled by the agent. The agent indicates status in a **Set** response message.
- **Traps** notification—The agent sends traps to notify the manager of significant events that occur on the network device.

The SNMP implementation in Junos OS contains:

- A master SNMP agent (known as the SNMP process or `snmpd`) that resides on the managed device and is managed by the NMS or host.
- Various subagents that reside on different modules of Junos OS, such as the Routing Engine, and are managed by the master SNMP agent (`snmpd`).



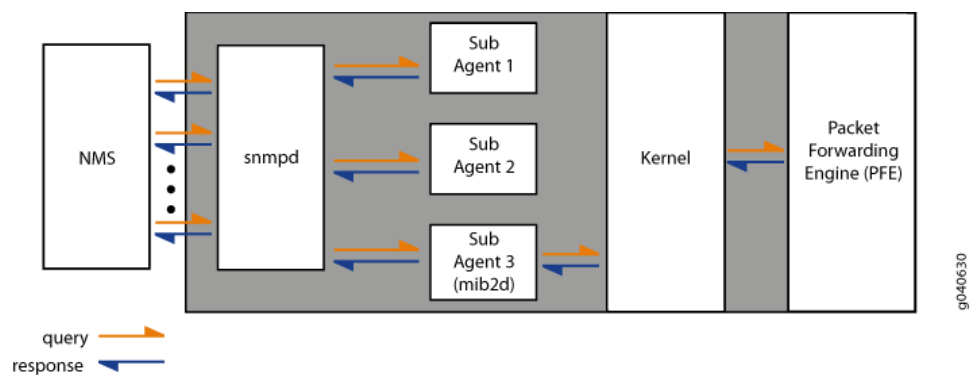
NOTE: By default, SNMP is not enabled on devices running Junos OS. For information about enabling SNMP on a device running the Junos OS, see [“Configuring SNMP on Devices Running Junos OS” on page 85](#).

The SNMP implementation in Junos OS uses both standard (developed by the IETF and documented in RFCs) and enterprise-specific (developed and supported by specific vendors) MIBs.

In Junos OS, the management data is maintained by the `snmpd` at one level (for example, `snmpVacmMIB` and `snmpUsmMIB`), and the subagents at the next level (for example, routing MIBs and RMON MIBs). However, there is another level of data that is maintained neither by the master agent nor by the subagents. In such cases, the data is maintained by the Junos OS processes that share the data with the subagents when polled for SNMP data. Interface-related MIBs and Firewall MIBs are good examples of data maintained by Junos OS processes.

When a network management system polls the master agent for data, the master agent immediately shares the data with the network management system if the requested data is available with the master agent or one of the subagents. However, if the requested data does not belong to those categories that are maintained by the master agent or the subagents, the subagent polls the Junos OS kernel or the process that maintains that data. On receiving the required data, the subagent passes the response back to the master agent, which in turn passes it to the NMS.

The following illustration shows the communication flow among the NMS, SNMP process (`snmpd`), SNMP subagents, and the Junos OS processes.



When a significant event, most often an error or a failure, occurs on a network device, the SNMP agent sends notifications to the SNMP manager. The SNMP implementation in Junos OS supports two types of notifications: traps and informs. *Traps* are unconfirmed notifications, whereas *informs* are confirmed notifications. Informs are supported only on devices that support SNMP version 3 (SNMPv3) configuration.

Junos OS supports trap queuing to ensure that traps are not lost because of temporary unavailability of routes. Two types of queues, *destination queues* and a *throttle queue*, are formed to ensure delivery of traps and to control the trap traffic.

Junos OS forms a destination queue when a trap to a particular destination is returned because the host is not reachable, and adds the subsequent traps to the same destination to the queue. Junos OS checks for availability of routes every 30 seconds and sends the traps from the destination queue in a round-robin fashion.

If the trap delivery fails, the trap is added back to the queue, and the delivery attempt counter and the next delivery attempt timer for the queue are reset. Subsequent attempts occur at progressive intervals of 1 minute, 2 minutes, 4 minutes, and 8 minutes. The maximum delay between the attempts is 8 minutes, and the maximum number of

attempts is 10. After 10 unsuccessful attempts, the destination queue and all the traps in the queue are deleted.

Junos OS also has a throttle mechanism to control the number of traps (throttle threshold; default value of 500 traps) sent during a particular time period (throttle interval; default of 5 seconds) and to ensure consistency in trap traffic, especially when large number of traps are generated because of interface status changes. The throttle interval period begins when the first trap arrives at the throttle. All traps within the trap threshold are processed, and the traps beyond the threshold limit are queued.

The maximum size of trap queues—that is, throttle queue and destination queue put together—is 40,000. However, on EX Series Ethernet Switches, the maximum size of the trap queue is 1,000. The maximum size of any one queue is 20,000 for devices other than EX Series Switches. On EX Series Switches, the maximum size of one queue is 500. When a trap is added to the throttle queue, or if the throttle queue has exceeded the maximum size, the trap is added back on top of the destination queue, and all subsequent attempts from the destination queue are stopped for a 30-second period, after which the destination queue restarts sending the traps.

Related Documentation

- [FAQ: SNMP Support on Junos OS](#)
- [Configuring SNMP on Devices Running Junos OS on page 85](#)
- [Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS on page 195](#)
- [Optimizing the Network Management System Configuration for the Best Results on page 81](#)
- [Configuring Options on Managed Devices for Better SNMP Response Time on page 83](#)
- [Managing Traps and Informs](#)
- [Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage](#)

SNMPv3 Overview

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX](#)

In contrast to SNMP version 1 (SNMPv1) and SNMP version 2 (SNMPv2), SNMP version 3 (SNMPv3) supports authentication and encryption. SNMPv3 uses the user-based security model (USM) for message security and the view-based access control model (VACM) for access control. USM specifies authentication and encryption. VACM specifies access-control rules.

USM uses the concept of a user for which security parameters (levels of security, authentication, privacy protocols, and keys) are configured for both the agent and the manager. Messages sent using USM are better protected than messages sent with community strings, where passwords are sent in the clear. With USM, messages exchanged between the manager and the agent can have data integrity checking and data origin authentication. USM protects against message delays and message replays by using time indicators and request IDs. Encryption is also available.

To complement the USM, SNMPv3 uses the VACM, a highly granular access-control model for SNMPv3 applications. Based on the concept of applying security policies to the name of the groups querying the agent, the agent decides whether the group is allowed to view or change specific MIB objects. VACM defines collections of data (called views), groups of data users, and access statements that define which views a particular group of users can use for reading, writing, or receiving traps.

Trap entries in SNMPv3 are created by configuring the notify, notify filter, target address, and target parameters. The **notify** statement specifies the type of notification (trap) and contains a single tag. The tag defines a set of target addresses to receive a trap. The notify filter defines access to a collection of trap object identifiers (OIDs). The target address defines a management application's address and other attributes to be used in sending notifications. Target parameters define the message processing and security parameters to be used in sending notifications to a particular management target.

To configure SNMPv3, perform the following tasks:

- [Creating SNMPv3 Users on page 121](#)
- [Configuring MIB Views on page 111](#)
- [Defining Access Privileges for an SNMP Group on page 126](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring SNMP Informs on page 143](#)

**Related
Documentation**

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

CHAPTER 4

SNMP MIBs and Traps Supported by Junos OS

- [Enterprise-Specific SNMP MIBs Supported by Junos OS on page 21](#)
- [Standard SNMP MIBs Supported by Junos OS on page 32](#)
- [Standard SNMP Traps Supported by Junos OS on page 50](#)
- [Enterprise-Specific SNMP Traps Supported by Junos OS on page 57](#)

Enterprise-Specific SNMP MIBs Supported by Junos OS

Supported Platforms [ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX](#)

Junos OS supports the enterprise-specific MIBs listed in [Table 6 on page 21](#). For information about enterprise-specific SNMP MIB objects, see the [SNMP MIB Explorer](#).

Table 6: Enterprise-specific MIBs supported by Junos OS

Enterprise-Specific MIB	Description	Platforms
AAA Objects MIB	Provides support for monitoring user authentication, authorization, and accounting through the RADIUS, LDAP, SecurID, and local authentication servers.	SRX Series and vSRX
Access Authentication Objects MIB	Provides support for monitoring firewall authentication, including data about the users trying to access firewall-protected resources and the firewall authentication service itself.	SRX Series and vSRX
Alarm MIB	Provides information about alarms from the router chassis.	All platforms

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
Analyzer MIB	Provides information about analyzer and remote analyzer related to port mirroring on the EX Series Ethernet Switches. Port mirroring is a method used on enterprise switches to monitor and analyze traffic on the network. When port mirroring is enabled, copies of all (or a sample set of) packets are forwarded from one port of the switch to another port on the same switch (analyzer) or on another switch (remote analyzer) where the packet can be analyzed and studied.	EX Series, QFabric system, and QFX Series
Antivirus Objects MIB	Provides information about the antivirus engine, antivirus scans, and antivirus scan-related traps.	SRX Series and vSRX
ATM Class-of-Service MIB	Provides support for ATM interfaces and virtual connections.	ACX Series, M Series, and T Series
ATM MIB	Provides support for monitoring Asynchronous Transfer Mode, version 2 (ATM2) virtual circuit (VC) class-of-service (CoS) configurations. It also provides CoS queue statistics for all VCs that have CoS configured.	M Series, SRX Series, T Series and vSRX
BGP4 V2 MIB	Provides support for monitoring BGP peer-received prefix counters. It is based upon similar objects in the MIB documented in Internet draft draft-ietf-idr-bgp4-mibv2-03.txt, <i>Definitions of Managed Objects for the Fourth Version of BGP (BGP-4), Second Version</i> .	All platforms
Bidirectional Forwarding Detection MIB	Provides support for monitoring Bidirectional Forwarding Detection (BFD) sessions.	All platforms
Chassis Cluster MIB	Provides information about objects that are used whenever the state of the control link interfaces or fabric link interfaces changes (up to down or down to up) in a chassis cluster deployment.	SRX Series and vSRX
Chassis Definitions for Router Model MIB	Contains the object identifiers (OIDs) that are used by the Chassis MIB to identify platform and chassis components. The Chassis MIB provides information that changes often, whereas the Chassis Definitions for Router Model MIB provides information that changes less often.	ACX Series, M Series, MX Series, PTX Series, QFX Series, SRX550, SRX1500, and T Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
Chassis MIBs	Provides support for environmental monitoring (power supply state, board voltages, fans, temperatures, and air flow) and inventory support for the chassis, System Control Board (SCB), System and Switch Board (SSB), Switching and Forwarding Module (SFM), Switch Fabric Board (SFB), Flexible PIC Concentrators (FPCs), and PICs.	All platforms
Class-of-Service MIB	Provides support for monitoring interface output queue statistics per interface and per forwarding class.	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric system, QFX Series, SRX Series, T Series, and vSRX
Configuration Management MIB	Provides notification for configuration changes as SNMP traps. Each trap contains the time at which the configuration change was committed, the name of the user who made the change, and the method by which the change was made. A history of the last 32 configuration changes is kept in <code>jnxCmChgEventTable</code> .	All platforms
Destination Class Usage MIB	Provides support for monitoring packet counts based on the ingress and egress points for traffic transiting your networks. Ingress points are identified by the input interface. Egress points are identified by destination prefixes grouped into one or more sets, known as destination classes. One counter is managed per interface per destination class, up to a maximum of 16 counters per interface.	EX Series, M Series, SRX Series, T Series, and vSRX
DHCP MIB	Provides SNMP support (get and trap) for DHCP local server and relay configurations. It also provides support for bindings and leases tables, and for statistics.	M Series, MX Series, and T Series
DHCPv6 MIB	Provides SNMP support (get and trap) for DHCPv6 local server and relay configurations. It also provides support for bindings and leases tables, and for statistics.	M Series, MX Series, and T Series
Digital Optical Monitoring MIB	Provides support for the SNMP Get request for statistics and SNMP Trap notifications for alarms.	EX Series, M Series, MX Series, PTX Series, and T Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
DNS Objects MIB	Provides support for monitoring DNS proxy queries, requests, responses, and failures.	SRX Series and vSRX
Dynamic Flow Capture MIB	Provides support for monitoring the operational status of dynamic flow capture (DFC) PICs.	M Series and T Series
Ethernet MAC MIB	Monitors media access control (MAC) statistics on Gigabit Ethernet intelligent queuing (IQ) interfaces. It collects MAC statistics; for example, inocets , inframes , outocets , and outframes on each source MAC address and virtual LAN (VLAN) ID for each Ethernet port.	EX Series, M Series, MX Series, QFX Series, SRX300, SRX320, SRX340, SRX550, SRX1500 and T Series
Event MIB	Defines a generic trap that can be generated using an op script or event policy. This MIB provides the ability to specify a system log string and raise a trap if that system log string is found.	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric system, QFX Series, SRX1500, SRX300, SRX320, SRX340, SRX550, and T Series
Experimental MIB	Contains object identifiers for experimental MIBs.	ACX Series, M series, MX Series, and T series
EX Series MAC Notification MIB	Contains Juniper Networks' implementation of enterprise-specific MIB for Ethernet Mac Stats for EX Series.	EX Series
EX Series SMI MIB	Contains the Structure of Management Information for Juniper Networks EX Series platforms.	EX Series
Firewall MIB	Provides support for monitoring firewall filter counters. Routers must have the Internet Processor II ASIC to perform firewall monitoring.	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric system, QFX Series, SRX300, SRX320, SRX340, SRX550, SRX1500 and T Series
Flow Collection Services MIB	Provides statistics on files, records, memory, FTP, and error states of a monitoring services interface. It also provides SNMP traps for unavailable destinations, unsuccessful file transfers, flow overloading, and memory overloading.	M Series and T Series
GRE Keepalive Monitoring MIB	Provides support for monitoring generic routing encapsulation (GRE) keepalive status. This MIB also provides an SNMP trap when GRE keepalive status changes.	SRX Series and vSRX instances

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
Host Resources MIB	Extends the hrStorageTable object, providing a measure of the usage of each file system on the router in percentage format. Previously, the objects in the hrStorageTable measured the usage in allocation units— hrStorageUsed and hrStorageAllocationUnits —only. Using the percentage measurement, you can more easily monitor and apply thresholds on usage.	ACX Series, EX Series, M Series, MX Series, QFX Series, SRX300, SRX320, SRX340, SRX550, SRX1500 and T Series
Interface Accounting Forwarding Class MIB	Extends the Juniper Enterprise Interface MIB and provides support for monitoring statistics data for interface accounting and IETF standardization.	M Series, MX Series, SRX Series, and vSRX
Interface MIB	Extends the standard ifTable (RFC 2863) with additional statistics and Juniper Networks enterprise-specific chassis information.	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric system, QFX Series, SRX300, SRX320, SRX340, SRX550, SRX1500 and T Series
IP Forward MIB	Extends the standard IP Forwarding Table MIB (RFC 4292) to include CIDR forwarding information.	All platforms
IPsec Generic Flow Monitoring Object MIB	Based on jnx-ipsec-monitor-mib , this MIB provides support for monitoring IPsec and IPsec VPN management objects.	SRX Series and vSRX
IPsec Monitoring MIB	Provides operational and statistical information related to the IPsec and IKE tunnels on Juniper Networks routers.	M Series, SRX Series, and T Series
IPsec VPN Objects MIB	Provides support for monitoring IPsec and IPsec VPN management objects for Juniper security product lines. This MIB is an extension of jnx-ipsec-flow-mon.mib .	SRX Series
IPv4 MIB	Provides additional Internet Protocol version 4 (IPv4) address information, supporting the assignment of identical IPv4 addresses to separate interfaces.	All platforms
IPv6 and ICMPv6 MIB	Provides IPv6 and Internet Control Message Protocol version 6 (ICMPv6) statistics.	M series, MX Series, PTX Series, SRX Series, T Series, and vSRX

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
L2ALD MIB	<p>Contains information about the Layer 2 Address Learning Daemon (L2ALD) and related traps, such as the routing instance MAC limit trap and the interface MAC limit trap. This MIB also provides VLAN information in the jnxL2aldVlanTable table for Enhanced Layer 2 Software (ELS) EX Series and QFX Series switches.</p> <p>NOTE: Non-ELS EX Series switches support the VLAN MIB (jnxExVlanTable table) for VLAN information instead of this MIB. See the SNMP MIB Explorer.</p>	EX Series, MX Series, QFX Series, and T Series
L2CP MIB	<p>Provides information about Layer 2 Control Protocols (L2CP) based features. Currently, Junos OS supports only the jnxDot1dStpPortRootProtectEnabled, jnxDot1dStpPortRootProtectState, and jnxPortRootProtectStateChangeTrap objects.</p>	MX Series
L2TP MIB	Provides information about Layer 2 Transport Protocol (L2TP) tunnels and sessions.	M Series, MX Series, and T Series
LDP MIB	Provides LDP statistics and defines LDP label-switched path (LSP) notifications. LDP traps support only IPv4 standards.	ACX Series, M Series, PTX Series, SRX Series, and T Series
License MIB	Extends SNMP support to licensing information, and introduces SNMP traps that alert users when the licenses are about to expire, expire, or when the total number of users exceeds the number specified in the license.	M Series, MX Series, SRX Series, and T Series
Logical Systems MIB	Extend SNMP support to logical systems security profile through various MIBs defined under jnxLsysSecurityProfile .	SRX Series
MIMSTP MIB	Provides information about MSTP instances (that is, routing instances of type Virtual Switch/Layer 2 control, also known as virtual contexts), MSTIs within the MSTP instance, and VLANs associated with the MSTI.	MX Series and T Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
MPLS LDP MIB	<p>Contains object definitions as described in RFC 3815, <i>Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)</i>.</p> <p>NOTE: Objects in the MPLS LDP MIB were supported in earlier releases of Junos OS as a proprietary LDP MIB (mib-ldpmib.txt). Because the branch used by the proprietary LDP (mib-ldpmib.txt) conflicts with RFC 3812, the proprietary LDP MIB (mib-ldpmib.txt) has been deprecated and replaced by the enterprise-specific MPLS LDP MIB (mib-jnx-mpls-ldp.txt).</p>	ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series
MPLS MIB	<p>Provides MPLS information and defines MPLS notifications.</p> <p>NOTE: To collect information about MPLS statistics on transit routers, use the enterprise-specific RSVP MIB (mib-jnx-rsvp.txt) instead of the enterprise-specific MPLS MIB (mib-jnx-mpls.txt).</p>	ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, and T Series
MVPN MIB	<p>Contains objects that enable SNMP manager to monitor MVPN connections on the provider edge routers. The enterprise-specific MVPN MIB is the Juniper Networks extension of the IETF standard MIBs defined in Internet draft draft-ietf-l3vpn-mvpn-mib-03.txt, <i>MPLS/BGP Layer 3 VPN Multicast Management Information Base</i>.</p>	All platforms
NAT Objects MIB	<p>Provides support for monitoring network address translation (NAT).</p>	EX Series and SRX Series
NAT Resources-Monitoring MIB	<p>Provides support for monitoring NAT pools usage and NAT rules. Notifications of usage of NAT resources are also provided by this MIB. This MIB is currently supported on the Multiservices PIC and Multiservices DPC on M Series and MX Series routers only.</p>	M Series and MX Series
OTN Interface Management MIB	<p>Defines objects for managing Optical Transport Network (OTN) interfaces on devices running Junos OS.</p>	M Series, MX series, PTX Series, and T Series
Packet Forwarding Engine MIB	<p>Provides notification statistics for Packet Forwarding Engines.</p>	ACX Series, EX Series, M Series, PTX Series, SRX Series, and T Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
Packet Mirror MIB	Enables you to capture and view packet mirroring-related information. This MIB is currently supported by Junos OS for MX Series routers only. Packet mirroring traps are an extension of the standard SNMP implementation and are only available to SNMPv3 users.	MX Series
PAE Extension MIB	Extends the standard IEEE802.1x PAE Extension MIB, and contains information for Static MAC Authentication.	EX Series
Passive Monitoring MIB	Performs traffic flow monitoring and lawful interception of packets transiting between two routers.	M Series and T Series
Ping MIB	Extends the standard Ping MIB control table (RFC 2925). Items in this MIB are created when entries are created in pingCtlTable of the Ping MIB. Each item is indexed exactly as it is in the Ping MIB.	ACX Series, EX Series, M Series, MX Series, QFX Series, SRX Series, and T Series
Policy Objects MIB	Provides support for monitoring the security policies that control the flow of traffic from one zone to another.	SRX Series
Power Supply Unit MIB	Enables monitoring and managing of the power supply on a device running Junos OS.	EX Series and QFabric system
PPP MIB	Provides SNMP support for PPP-related information such as the type of authentication used, interface characteristics, status, and statistics. This MIB is supported on Common Edge PPP process, jpppd.	M Series and MX Series
PPPoE MIB	Provides SNMP support for PPPoE-related information such as the type of authentication used, interface characteristics, status, and statistics. This MIB is supported on Common Edge PPPoE process, jpppoed.	M Series and MX Series
Pseudowire ATM MIB	Extends the standard Pseudowire MIB, and defines objects used for managing the ATM pseudowires in Juniper products. The enterprise-specific Pseudowire ATM MIB is the Juniper Networks implementation of RFC 5605, <i>Managed Objects for ATM over Packet Switched Networks (PSNs)</i> .	M Series and MX Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
Pseudowire TDM MIB	Extends the standard Pseudowire MIB, and contains information about configuration and statistics for specific pseudowire types. The enterprise-specific Pseudowire TDM MIB is the Juniper Networks implementation of the standard Managed Objects for TDM over Packet Switched Network MIB (draft-ietf-pwe3-tdm-mib-08.txt).	ACX Series, M Series, and T Series
PTP MIB	Monitors the operation of PTP clocks within the network.	MX Series
Real-Time Performance Monitoring MIB	Provides real-time performance-related data and enables you to access jitter measurements and calculations using SNMP.	EX Series, M Series, MX Series, SRX Series, and T Series
Reverse-Path-Forwarding MIB	Monitors statistics for traffic that is rejected because of reverse-path-forwarding (RPF) processing.	All platforms
RMON Events and Alarms MIB	Supports the Junos OS extensions to the standard Remote Monitoring (RMON) Events and Alarms MIB (RFC 2819). The extension augments alarmTable with additional information about each alarm. Two new traps are also defined to indicate when problems are encountered with an alarm.	All platforms
RSVP MIB	Provides information about RSVP-traffic engineering sessions that correspond to MPLS LSPs on transit routers in the service provider core network. NOTE: To collect information about MPLS statistics on transit routers, use the enterprise-specific RSVP MIB (mib-jnx-rsvp.txt) instead of the enterprise-specific MPLS MIB (mib-jnx-mpls.txt).	ACX Series, M Series, MX Series, PTX Series, and T Series
Security Interface Extension Objects MIB	Provides support for the security management of interfaces.	EX Series, SRX Series, and vSRX
Security Screening Objects MIB	Defines the MIB for the Juniper Networks Enterprise Firewall screen functionality.	SRX Series and vSRX
Services PIC MIB	Provides statistics for Adaptive Services (AS) PICs and defines notifications for AS PICs.	M Series and T Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
SNMP IDP MIB	Contains Juniper Networks' implementation of enterprise specific MIB for IDP.	SRX Series and vSRX
SONET APS MIB	Monitors any SONET interface that participates in Automatic Protection Switching (APS).	M Series and T Series
SONET/SDH Interface Management MIB	Monitors the current alarm for each SONET/SDH interface.	M Series and T Series
Source Class Usage MIB	Counts packets sent to customers by performing a lookup on the IP source address and the IP destination address. The Source Class Usage (SCU) MIB makes it possible to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge.	M Series, T Series, and SRX Series
SPU Monitoring MIB	Provides support for monitoring SPUs on SRX5600 and SRX5800 devices.	SRX Series and vSRX
Structure of Management Information MIB	Explains how the Juniper Networks enterprise-specific MIBs are structured.	ACX Series, EX Series, M Series, MX series, QFX Series, SRX Series, T Series and vSRX
Structure of Management Information MIB for EX Series Ethernet Switches	Defines a MIB branch for switching-related MIB definitions for the EX Series Ethernet Switches.	EX Series
Structure of Management Information MIB for SRX Series	Contains object identifiers (OIDs) for the security branch of the MIBs used in Junos OS for SRX Series devices, services, and traps.	SRX Series and vSRX
Subscriber MIB	Provides SNMP support for subscriber-related information.	ACX Series, MX Series, and T Series
System Log MIB	Enables notification of an SNMP trap-based application when an important system log message occurs.	EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, and T Series
Traceroute MIB	Supports the Junos OS extensions of traceroute and remote operations. Items in this MIB are created when entries are created in the <code>traceRouteCtlTable</code> of the Traceroute MIB. Each item is indexed exactly the same way as it is in the Traceroute MIB.	EX Series, M Series, MX Series, SRX Series, T Series, and vSRX

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
Tunnel Stats MIB	Supports monitoring of tunnel statistics for IPV4 over IPV6 tunnels. This MIB currently displays three counters: tunnel count in rpd, tunnel count in Kernel, and tunnel count in the Packet Forwarding Engine.	all platforms
Utility MIB	Provides SNMP support for exposing the Junos OS data and has tables that contain information about each type of data, such as integer and string.	EX Series, M Series, MX Series, QFabric system, QFX Series, SRX Series, T Series, and vSRX
Virtual Chassis MIB	Contains information about the virtual chassis on the EX Series Ethernet Switches and the MX Series.	EX Series and MX Series
VLAN MIB	<p>Contains information about prestandard IEEE 802.10 VLANs and their association with LAN emulation clients.</p> <p>NOTE: For ELS EX Series switches and QFX Series switches, VLAN information is provided in the L2ALD MIB in the jnxL2aldVlanTable table instead of in this MIB. See the SNMP MIB Explorer for details.</p> <p>Non-ELS EX Series Ethernet switches use the jnxExVlanTable table in this MIB to provide VLAN configuration information, and the jnxVlanTable table in this MIB has been deprecated and is no longer used.</p>	EX Series and QFX Series
VPLS MIBs	<p>Provides information about generic, BGP-based, and LDP-based VPLS, and pseudowires associated with the VPLS networks. The enterprise-specific VPLS MIBs are Juniper Networks extensions of the following IETF standard MIBs defined in Internet draft draft-ietf-l2vpn-vpls-mib-05.txt, and are implemented as part of the jnxExperiment branch:</p> <ul style="list-style-type: none"> • VPLS-Generic-Draft-01-MIB implemented as mib-jnx-vpls-generic.txt • VPLS-BGP-Draft-01-MIB implemented as mib-jnx-vpls-bgp.txt • VPLS-LDP-Draft-01-MIB implemented as mib-jnx-vpls-ldp.txt 	M Series, MX Series, and T Series

Table 6: Enterprise-specific MIBs supported by Junos OS (*continued*)

Enterprise-Specific MIB	Description	Platforms
VPN Certificate Objects MIB	Provides support for monitoring the local and CA certificates loaded on the router.	EX Series, SRX Series, and vSRX
VPN MIB	Provides monitoring for Layer 3 VPNs, Layer 2 VPNs, and virtual private LAN service (VPLS) (read access only).	ACX Series, EX Series, M Series, MX Series, and T Series

For information about enterprise-specific SNMP MIB objects, see the [SNMP MIB Explorer](#).

Related Documentation

- *Network Management Administration Guide*
- [Standard SNMP MIBs Supported by Junos OS on page 32](#)
- [Enterprise-Specific SNMP Traps Supported by Junos OS on page 57](#)

Standard SNMP MIBs Supported by Junos OS

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Junos OS supports the Standard MIBs listed in [Table 7 on page 32](#).



NOTE: For details on SNMP MIB support on EX4600 switches, QFX Series switches, and QFabric systems, see *SNMP MIBs Support*.

Table 7: Standard MIBs supported by Junos OS

Standard MIB	Exceptions	Platforms
IEEE 802.1ab section 12.1, <i>Link Layer Discovery Protocol (LLDP) MIB</i>	EX Series implementation of LLDP MIB supports both IPv4 and IPv6 configuration.	EX Series and MX Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
IEEE, 802.3ad, <i>Aggregation of Multiple Link Segments</i>	<p>Supported tables and objects:</p> <ul style="list-style-type: none"> • dot3adAggPortTable, dot3adAggPortListTable, dot3adAggTable, and dot3adAggPortStatsTable <p>NOTE: EX Series switches do not support the dot3adAggPortTable and dot3adAggPortStatsTable.</p> <ul style="list-style-type: none"> • dot3adAggPortDebugTable (only dot3adAggPortDebugRxState, dot3adAggPortDebugMuxState, dot3adAggPortDebugActorSyncTransitionCount, dot3adAggPortDebugPartnerSyncTransitionCount, dot3adAggPortDebugActorChangeCount, and dot3adAggPortDebugPartnerChangeCount) <p>NOTE: EX Series switches do not support the dot3adAggPortDebugTable.</p> <ul style="list-style-type: none"> • dot3adTablesLastChanged 	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, and vSRX

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
IEEE, 802.1ag, <i>Connectivity Fault Management</i>	Supported tables and objects: <ul style="list-style-type: none"> • dot1agCfmMdTableNextIndex • dot1agCfmMdTable (except dot1agCfmMdMhfdPermission) • dot1agCfmMaNetTable • dot1agCfmMaMepListTable • dot1agCfmDefaultMdDefLevel • dot1agCfmDefaultMdDefMhfCreation • dot1agCfmMepTable (except dot1agCfmMepLbrBadMsdu, dot1agCfmMepTransmitLbmVlanPriority, dot1agCfmMepTransmitLbmVlanDropEnable, dot1agCfmMepTransmitLtmFlags, dot1agCfmMepPbbTeCanReportPbbTePresence, dot1agCfmMepPbbTeTrafficMismatchDefect, dot1agCfmMepPbbTeTransmitLbmLtmReverseVid, dot1agCfmMepPbbTeMismatchAlarm, dot1agCfmMepPbbTeLocalMismatchDefect, and dot1agCfmMepPbbTeMismatchSinceReset) • dot1agCfmLtrTable (except dot1agCfmLtrChassisIdSubtype, dot1agCfmLtrChassisId, dot1agCfmLtrManAddressDomain, dot1agCfmLtrManAddress, dot1agCfmLtrIngressPortIdSubtype, dot1agCfmLtrIngressPortId, dot1agCfmLtrEgressPortIdSubtype, dot1agCfmLtrEgressPortId, and dot1agCfmLtrOrganizationSpecificTlv) • dot1agCfmMepDbTable (except dot1agCfmMebDbChassisIdSubtype, dot1agCfmMebDbChassisId, dot1agCfmMebDbManAddressDomain, and dot1agCfmMebDbManAddress) 	EX Series, MX Series, and QFX Series
IEEE, 802.1ap, <i>Management Information Base (MIB) definitions for VLAN Bridges</i>	Supported tables and objects: <ul style="list-style-type: none"> • ieee8021CfmStackTable • ieee8021CfmVlanTable • ieee8021CfmDefaultMdTable (except ieee8021CfmDefaultMdIdPermission) • ieee8021CfmMaCompTable (except ieee8021CfmMaCompldPermission) 	MX Series
RFC 1155, <i>Structure and Identification of Management Information for TCP/IP-based Internets</i>	No exceptions	All platforms
RFC 1157, <i>A Simple Network Management Protocol (SNMP)</i>	No exceptions	All platforms

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 1195, <i>Use of OSI IS-IS for Routing in TCP/IP and Dual Environments</i>	Supported tables and objects: <ul style="list-style-type: none"> • isisSystem • isisMANAreaAddr • isisAreaAddr • isisSysProtSupp • isisSummAddr • isisCirc • isisCircLevel • isisPacketCount • isisISAdj • isisISAdjAreaAddr • isisAdjIPAddr • isisISAdjProtSupp • isisRa • isisIPRA 	All platforms
RFC 1212, <i>Concise MIB Definitions</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 1213, <i>Management Information Base for Network Management of TCP/IP-Based Internets: MIB-II</i>	Junos OS supports the following areas: <ul style="list-style-type: none"> • MIB II and its SNMP version 2 derivatives, including: <ul style="list-style-type: none"> • Statistics counters • IP, except for ipRouteTable, which has been replaced by ipCidrRouteTable (RFC 2096, <i>IP Forwarding Table MIB</i>) • SNMP management • Interface management • SNMPv1 Get, GetNext requests, and version 2 GetBulk request • Junos OS-specific secured access list • Master configuration keywords • Reconfigurations upon SIGHUP 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 1215, <i>A Convention for Defining Traps for use with the SNMP</i>	Junos OS supports only MIB II SNMP version 1 traps and version 2 notifications.	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 1406, <i>Definitions of Managed Objects for the DS1 and E1 Interface Types</i>	Junos OS supports T1 MIB.	ACX Series, M Series, SRX Series, and T Series
RFC 1407, <i>Definitions of Managed Objects for the DS3/E3 Interface Type</i>	Junos OS supports T3 MIB.	M Series and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 1471, <i>Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol</i>	Supported tables and objects: <ul style="list-style-type: none"> • pppLcp 1 object • pppLinkStatustable table • pppLinkConfigTable table 	M Series, MX Series, and PTX Series
RFC 1657, <i>Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, and T Series
RFC 1695, <i>Definitions of Managed Objects for ATM Management Version 8.0 Using SMIv2</i>	No exceptions	ACX Series, M Series, PTX Series, and T Series
RFC 1850, <i>OSPF Version 2 Management Information Base</i>	Unsupported tables, objects, and traps: <ul style="list-style-type: none"> • ospfOriginateNewLsas object • ospfRxNewLsas object • The host table • ospfOriginateLSA trap • ospfLsdbOverflow trap • ospfLsdbApproachingOverflow trap 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 1901, <i>Introduction to Community-based SNMPv2</i>	No exceptions	All platforms
RFC 2011, <i>SNMPv2 Management Information Base for the Internet Protocol Using SMIv2</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, and T Series
RFC 2012, <i>SNMPv2 Management Information Base for the Transmission Control Protocol Using SMIv2</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2013, <i>SNMPv2 Management Information Base for the User Datagram Protocol Using SMIv2</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 2024, <i>Definitions of Managed Objects for Data Link Switching Using SMIv2</i>	Unsupported tables, objects, and traps with read-only access: <ul style="list-style-type: none"> • dlswInterface object group • dlswSdlc object group • dlswDirLocateMacTable table • dlswDirNBTable table • dlswDirLocateNBTable table • dlswCircuitDiscReasonLocal tabular object • dlswCircuitDiscReasonRemote tabular object • dlswDirMacCacheNextIndex scalar object • dlswDirNBCacheNextIndex scalar object 	M Series, MX Series, and T Series
RFC 2096, <i>IP Forwarding Table MIB</i> NOTE: RFC 2096 has been replaced by RFC 4292. However, Junos OS currently supports both RFC 2096 and RFC 4292.	The ipCidrRouteTable has been extended to include the tunnel name when the next hop is through an RSVP-signaled LSP.	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2115, <i>Management Information Base for Frame Relay DTEs Using SMIv2</i>	Unsupported table and objects: <ul style="list-style-type: none"> • frCircuitTable • frErrTable 	M Series, MX Series, SRX Series, and T Series
RFC 2233, <i>The Interfaces Group MIB Using SMIv2</i> NOTE: RFC 2233 has been replaced by RFC 2863, IF MIB. However, Junos OS supports both RFC 2233 and RFC 2863.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2287, <i>Definitions of System-Level Managed Objects for Applications</i>	Supported tables and objects: <ul style="list-style-type: none"> • sysApplInstallPkgTable • sysApplInstallElmtTable • sysApplElmtRunTable • sysApplMapTable 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2465, <i>Management Information Base for IP Version 6: Textual Conventions and General Group</i> (except for IPv6 interface statistics)	Junos OS does not support IPv6 interface statistics.	ACX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 2495, <i>Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types</i>	Unsupported tables, objects, and traps: <ul style="list-style-type: none"> • dsx1FarEndConfigTable • dsx1FarEndCurrentTable • dsx1FarEndIntervalTable • dsx1FarEndTotalTable • dsx1FracTable 	ACX Series, M Series, SRX Series, and T Series
RFC 2515, <i>Definitions of Managed Objects for ATM Management</i>	Unsupported table and objects: <ul style="list-style-type: none"> • atmVpCrossConnectTable • atmVcCrossConnectTable • aal5VccTable 	ACX Series, M Series, and T Series
RFC 2570, <i>Introduction to Version 3 of the Internet-standard Network Management Framework</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2571, <i>An Architecture for Describing SNMP Management Frameworks</i> (read-only access) NOTE: RFC 2571 has been replaced by RFC 3411. However, Junos OS supports both RFC 2571 and RFC 3411.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2572, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> (read-only access) NOTE: RFC 2572 has been replaced by RFC 3412. However, Junos OS supports both RFC 2572 and RFC 3412.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2576, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i> NOTE: RFC 2576 has been replaced by RFC 3584. However, Junos OS supports both RFC 2576 and RFC 3584.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2578, <i>Structure of Management Information Version 2 (SMIv2)</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2579, <i>Textual Conventions for SMIv2</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 2580, <i>Conformance Statements for SMIPv2</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2662, <i>Definitions of Managed Objects for ADSL Lines</i>	No exceptions	M Series, MX Series, SRX Series, and T Series
RFC 2665, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i> NOTE: The list of managed objects specified in RFC 2665 has been updated by RFC 3635 by including information useful for the management of 10-Gigabit per second Ethernet interfaces.	For M Series, T Series, and MX Series, the SNMP counters do not count the Ethernet header and frame check sequence (FCS). Therefore, the Ethernet header bytes and the FCS bytes are not included in the following four tables: <ul style="list-style-type: none"> • ifInOctets • ifOutOctets • ifHCInOctets • ifHCOctets However, the EX switches adhere to RFC 2665.	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	Unsupported table and objects: <ul style="list-style-type: none"> • vrrpStatsPacketLengthErrors NOTE: Junos OS does not support this standard for row creation and the Set operation.	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2790, <i>Host Resources MIB</i>	Supported tables and objects: <ul style="list-style-type: none"> • hrStorageTable NOTE: The file systems / , /config , /var , and /tmp always return the same index number. When SNMP restarts, the index numbers for the remaining file systems might change. <ul style="list-style-type: none"> • hrSystem group • hrSWInstalled group • hrProcessorTable 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 2819, <i>Remote Network Monitoring Management Information Base</i>	Supported tables and objects: <ul style="list-style-type: none"> • etherStatsTable (for Ethernet interfaces only), alarmTable, eventTable, and logTable are supported on all devices running Junos OS. • historyControlTable and etherHistoryTable (except etherHistoryUtilization object) are supported only on EX Series switches. 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2863, <i>The Interfaces Group MIB</i> NOTE: RFC 2863 replaces RFC 2233. However, Junos OS supports both RFC 2233 and RFC 2863.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2864, <i>The Inverted Stack Table Extension to the Interfaces Group MIB</i>	No exceptions	M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2922, <i>The Physical Topology (PTOPO) MIB</i>	Supported objects: <ul style="list-style-type: none"> • ptopoConnDiscAlgorithm • ptopoConnAgentNetAddrType • ptopoConnAgentNetAddr • ptopoConnMultiMacSASeen • ptopoConnMultiNetSASeen • ptopoConnIsStatic • ptopoConnLastVerifyTime • ptopoConnRowStatus 	EX Series and SRX Series
RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	Supported tables and objects: <ul style="list-style-type: none"> • pingCtlTable • pingResultsTable • pingProbeHistoryTable • pingMaxConcurrentRequests • traceRouteCtlTable • traceRouteResultsTable • traceRouteProbeHistoryTable • traceRouteHopsTable 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2932, <i>IPv4 Multicast Routing MIB</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 2934, <i>Protocol Independent Multicast MIB for IPv4</i> NOTE: In Junos OS, RFC 2934 is implemented based on a draft version, <i>pimmib.mib</i> , of the now standard RFC.	Support for the pimNeighborLoss trap was added in Release 11.4.	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 2981, <i>Event MIB</i>	No exceptions	ACX Series, M Series, MX Series, PTX Series, and T Series
RFC 3014, <i>Notification Log MIB</i>	No exceptions	ACX Series, M Series, MX Series, PTX Series, and T Series
RFC 3019, <i>IP Version 6 Management Information Base for The Multicast Listener Discovery Protocol</i>	No exceptions	M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3410, <i>Introduction and Applicability Statements for Internet-Standard Management Framework</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> NOTE: RFC 3411 replaces RFC 2571. However, Junos OS supports both RFC 3411 and RFC 2571.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3412, <i>Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)</i> NOTE: RFC 3412 replaces RFC 2572. However, Junos OS supports both RFC 3412 and RFC 2572.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3413, <i>Simple Network Management Protocol (SNMP) Applications</i>	Unsupported tables and objects: <ul style="list-style-type: none"> Proxy MIB 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3414, <i>User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3415, <i>View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 3416, <i>Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)</i> NOTE: RFC 3416 replaces RFC 1905, which was supported in earlier versions of Junos OS.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3417, <i>Transport Mappings for the Simple Network Management Protocol (SNMP)</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3418, <i>Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)</i> NOTE: RFC 3418 replaces RFC 1907, which was supported in earlier versions of Junos OS.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3498, <i>Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures</i> (implemented under the Juniper Networks enterprise branch [jnxExperiment])	No exceptions	M Series and T Series
RFC 3584, <i>Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3591 <i>Managed Objects for the Optical Interface Type</i>	Supported tables and objects: <ul style="list-style-type: none"> • <code>optIfOTMnTable</code> (except <code>optIfOTMnOpticalReach</code>, <code>optIfOTMnInterfaceType</code>, and <code>optIfOTMnOrder</code>) • <code>optIfOChConfigTable</code> (except <code>optIfOChDirectionality</code> and <code>optIfOChCurrentStatus</code>) • <code>optIfOTUkConfigTable</code> (except <code>optIfOTUkTraceIdentifierAccepted</code>, <code>optIfOTUkTIMDetMode</code>, <code>optIfOTUkTIMActEnabled</code>, <code>optIfOTUkTraceIdentifierTransmitted</code>, <code>optIfOTUkDEGThr</code>, <code>optIfOTUkDEGM</code>, <code>optIfOTUkSinkAdaptActive</code>, and <code>optIfOTUkSourceAdaptActive</code>) • <code>optIfODUkConfigTable</code> (except <code>optIfODUkPositionSeqCurrentSize</code> and <code>optIfODUkTtpPresent</code>) 	M Series, MX Series, PTX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 3592, <i>Definitions of Managed Objects for the Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Interface Type</i>	No exceptions	M Series, MX Series, and T Series
RFC 3621, <i>Power Ethernet MIB</i>	No exceptions	EX Series
RFC 3635, <i>Definitions of Managed Objects for the Ethernet-like Interface Types</i>	Unsupported tables and objects: <ul style="list-style-type: none"> • dot3StatsRateControlAbility • dot3StatsRateControlStatus in dot3StatsEntry table <p>NOTE: The values of the following objects in dot3HCStatsEntry table will be always zero for both 32-bit counters and 64-bit counters:</p> <ul style="list-style-type: none"> • dot3HCStatsSymbolErrors • dotHCStatsInternalMacTransmitErrors 	MX Series
RFC 3637, <i>Definitions of Managed Objects for the Ethernet WAN Interface Sublayer</i>	Unsupported tables and objects: <ul style="list-style-type: none"> • etherWisDeviceTable, • etherWisSectionCurrentTable • etherWisFarEndPathCurrentTable 	M Series, MX Series, PTX Series, and T Series
RFC 3811, <i>Definitions of Textual Conventions (TCs) for Multiprotocol Label Switching (MPLS) Management</i>	No exceptions	ACX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 3812, <i>Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)</i> (read-only access)	<p>MPLS tunnels as interfaces are not supported.</p> <p>mplsTunnelCHopTable is supported on ingress routers only.</p> <p>NOTE: The branch used by the proprietary LDP MIB (ldpmib.mib) conflicts with RFC 3812. ldpmib.mib has been deprecated and replaced by jnx-mpls-ldp.mib.</p> <p>Unsupported tables and objects:</p> <ul style="list-style-type: none"> • mplsTunnelResourceMeanRate in TunnelResource table • mplsTunnelResourceMaxBurstSize in TunnelResource table • mplsTunnelResourceMeanBurstSize in TunnelResource table • mplsTunnelResourceExBurstSize in TunnelResource table • mplsTunnelResourceWeight in TunnelResource table • mplsTunnelPerfTable • mplsTunnelCRLDResTable 	ACX Series, M Series, MX Series, PTX Series, and T Series
RFC 3813, <i>Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)</i>	<p>Unsupported tables and objects (read-only access):</p> <ul style="list-style-type: none"> • mplsInterfacePerfTable • mplsInSegmentPerfTable • mplsOutSegmentPerfTable • mplsInSegmentMapTable • mplsXCUp • mplsXCDown 	ACX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3826, <i>The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 3877, <i>Alarm Management Information Base</i>	<ul style="list-style-type: none"> • Junos OS does not support the alarmActiveStatsTable. • Traps that do not conform to the alarm model are not supported. However, these traps can be redefined to conform to the alarm model. 	MX Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 3896, <i>Definitions of Managed Objects for the DS3/E3 Interface Type</i>	Unsupported tables and objects: <ul style="list-style-type: none"> • dsx3FarEndConfigTable • dsx3FarEndCurrentTable • dsx3FarEndIntervalTable • dsx3FarEndTotalTable • dsx3FracTable 	M Series and T Series
RFC 4087, <i>IP Tunnel MIB</i>	Describes MIB objects in the following tables for managing tunnels of any type over IPv4 and IPv6 networks: <ul style="list-style-type: none"> • tunnelIfTable—Provides information about the tunnels known to a router. • tunnelInetConfigTable—Assists dynamic creation of tunnels and provides mapping from end-point addresses to the current interface index value. <p>NOTE: Junos OS supports MAX-ACCESS of read-only for all the MIB objects in tunnelIfTable and tunnelInetConfigTable tables.</p>	M Series, MX Series, and T Series
RFC 4133, <i>Entity MIB</i>	Unsupported tables and objects: <ul style="list-style-type: none"> • entityLogicalGroup table • entPhysicalMfgDate and entPhysicalUris objects in entityPhysical2Group table • entLPMappingTable and entPhysicalContainsTable in entityMappingGroup table • entityNotificationsGroup table 	Only MX240, MX480, and MX960 routers, and EX2200 and EX3300 switches
RFC 4188, <i>Definitions of Managed Objects for Bridges</i>	<ul style="list-style-type: none"> • Supports 802.1D STP(1998) • Supported subtrees and objects: <ul style="list-style-type: none"> • dot1dStp subtree is supported on MX Series 3D Universal Edge Routers. • dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus objects from the dot1dTpFdbTable of the dot1dTp subtree are supported on EX Series Ethernet Switches. • dot1dTpLearnedEntryDiscards and dot1dTpAgingTime objects are supported on M Series and T Series routers. 	MX Series, EX Series, and M Series and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 4268, <i>Entity State MIB</i>	No exceptions	Only MX240, MX480, and MX960 routers, and EX2200 and EX3300 switches
RFC 4273, <i>Definitions of Managed Objects for BGP-4</i>	Supported tables and objects: <ul style="list-style-type: none"> • <code>jnxBgpM2PrefixInPrefixesAccepted</code> • <code>jnxBgpM2PrefixInPrefixesRejected</code> 	ACX Series, EX Series, M Series, MX Series, SRX Series, and T Series
RFC 4292, <i>IP Forwarding MIB</i>	Supported tables and objects: <ul style="list-style-type: none"> • <code>inetCidrRouteTable</code> • <code>inetCidrRouteNumber</code> • <code>inetCidrRouteDiscards</code> <p>NOTE: Junos OS currently supports these MIB objects that will be deprecated in future releases: <code>ipCidrRouteTable</code>, <code>ipCidrRouteNumber</code>, and <code>ipCidrRouteDiscards</code>.</p>	ACX Series, EX Series, M Series, MX Series, PTX Series, and T Series
RFC 4293, <i>Management Information Base for the Internet Protocol (IP)</i>	Supports only the mandatory groups.	MX Series and EX Series
RFC 4318, <i>Definitions of Managed Objects for Bridges with Rapid Spanning Tree Protocol</i>	Supports 802.1w and 802.1t extensions for RSTP.	EX Series, M Series, MX Series, and T Series
RFC 4363b, <i>Q-Bridge VLAN MIB</i>	No exceptions	MX Series and EX Series
RFC 4382, <i>MPLS/BGP Layer 3 Virtual Private Network (VPN) MIB</i>	Supported tables and objects: <ul style="list-style-type: none"> • <code>mplsL3VpnActiveVrfs</code> • <code>mplsL3VpnConfiguredVrfs</code> • <code>mplsL3VpnConnectedInterfaces</code> • <code>mplsL3VpnVrfConfMidRteThresh</code> • <code>mplsL3VpnVrfConfHighRteThresh</code> • <code>mplsL3VpnIfConfRowStatus</code> • <code>mplsL3VpnIfL3RcvThresh</code> • <code>mplsL3VpnNotificationEnable</code> • <code>mplsL3VpnVrfConfMaxPossRts</code> • <code>mplsL3VpnVrfConfRteMxThreshTime</code> • <code>mplsL3VpnVrfOperStatus</code> • <code>mplsL3VpnVrfPerfCurrNumRoutes</code> • <code>mplsL3VpnVrfPerfTable</code> • <code>mplsL3VpnVrfRteTable</code> • <code>mplsVpnVrfRTTable</code> • <code>mplsL3VpnVrfTable</code> • <code>mplsL3VpnIfConfTable</code> 	EX Series, M Series, MX Series, PTX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 4444, <i>IS-IS MIB</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 4668, <i>RADIUS Accounting Client Management Information Base (MIB) for IPv6</i> (read-only access)	No exceptions	MX Series
RFC 4670, <i>RADIUS Accounting Client Management Information Base (MIB)</i> (read-only access)	No exceptions	MX Series
RFC 4801, <i>Definitions of Textual Conventions for Generalized Multiprotocol Label Switching (GMPLS) Management Information Base (MIB)</i> (read-only access)	No exceptions	M Series, MX Series, and T Series
RFC 4802, <i>Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering (TE) Management Information Base (MIB)</i> (read-only access)	Unsupported tables and objects: <ul style="list-style-type: none"> • gmplsTunnelReversePerfTable • gmplsTeScalars • gmplsTunnelTable • gmplsTunnelARHopTable • gmplsTunnelCHopTable • gmplsTunnelErrorTable 	M Series, MX Series, and T Series
RFC 4803, <i>Generalized Multiprotocol Label Switching (GMPLS) Label Switching Router (LSR) Management Information Base (MIB)</i> (read-only access)	Unsupported tables and objects: <ul style="list-style-type: none"> • gmplsLabelTable • gmplsOutsegmentTable 	M Series, MX Series, and T Series
NOTE: The tables in GMPLS TE (RFC 4802) and LSR (RFC 4803) MIBs are extensions of the corresponding tables from the MPLS TE (RFC 3812) and LSR (RFC 3813) MIBs and use the same index as the MPLS MIB tables.		
RFC 5132, <i>IP Multicast MIB</i>	Unsupported table: <ul style="list-style-type: none"> • ipMcastZoneTable 	All platforms
NOTE: This RFC obsoletes RFC2932.		

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
RFC 5643, <i>Management Information Base for OSPFv3</i> (read-only access)	Unsupported tables and objects: <ul style="list-style-type: none"> • ospfv3HostTable • ospfv3CfgrNbrTable • ospfv3ExitOverflowInterval • ospfv3ReferenceBandwidth • ospfv3RestartSupport • ospfv3RestartInterval • ospfv3RestartStrictLsaChecking • ospfv3RestartStatus • ospfv3RestartAge • ospfv3RestartExitReason • ospfv3NotificationEnable • ospfv3StubRouterSupport • ospfv3StubRouterAdvertisement • ospfv3DiscontinuityTime • ospfv3RestartTime • ospfv3AreaNssaTranslatorRole • ospfv3AreaNssaTranslatorState • ospfv3AreaNssaTranslatorStabInterval • ospfv3AreaNssaTranslatorEvents • ospfv3AreaTEEnabled • ospfv3IfMetricValue • ospfv3IfDemandNbrProbe 	M Series, MX Series, PTX Series, SRX Series, and T Series
RFC 6527, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol Version 3 (VRRPv3)</i>	<ul style="list-style-type: none"> • Row creation • The Set operation • Unsupported tables and objects: <ul style="list-style-type: none"> • vrrpv3StatisticsRowDiscontinuityTime • vrrpv3StatisticsPacketLengthErrors 	ACX Series
ESO Consortium MIB, which can be found at http://www.snmp.com/eso/ NOTE: The ESO Consortium MIB has been replaced by RFC 3826.	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
Internet Assigned Numbers Authority, <i>IANAiftype Textual Convention MIB</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
Internet draft draft-ietf-atommib-sonetaps-mib-10.txt, <i>Definitions of Managed Objects for SONET Linear APS Architectures</i>	As defined under the Juniper Networks enterprise branch [jnxExperiment] only	M Series, MX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
Internet draft draft-ietf-bfd-mib-02.txt, <i>Bidirectional Forwarding Detection Management Information Base</i>	(Represented by mib-jnx-bfd-exp.txt and implemented under the Juniper Networks enterprise branch [jnxExperiment]. Read only. Includes bfdSessUp and bfdSessDown traps. Does not support bfdSessPerfTable and bfdSessMapTable .)	ACX Series, EX Series, M Series, MX Series, SRX Series, and T Series
Internet draft draft-ietf-idmr-igmp-mib-13.txt, <i>Internet Group Management Protocol (IGMP) MIB</i>	No exceptions	EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
Internet draft draft-ietf-idmr-pim-mib-09.txt, <i>Protocol Independent Multicast (PIM) MIB</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
Internet draft draft-ietf-isis-wg-mib-07.txt, <i>Management Information Base for IS-IS</i> NOTE: Replaced with RFC 4444, <i>IS-IS MIB</i> in Junos OS Release 11.3 and later.	Unsupported tables and objects: <ul style="list-style-type: none"> • isisISAdjTable • isisISAdjAreaAddrTable • isisISAdjIPAddrTable • isisISAdjProtSuppTable 	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series
Internet draft draft-ietf-l3vpn-mvpn-mib-03.txt, <i>MPLS/BGP Layer 3 VPN Multicast Management Information Base</i>	(Implemented under the Juniper Networks enterprise branch [jnxExperiment]. OID for jnxMvpnExperiment is .1.3.6.1.4.1.2636.5.12. Read only. Includes jnxMvpnNotifications traps.)	M Series, MX Series, and T Series
Internet draft draft-ietf-mppls-mldp-mib-02.txt, <i>Definitions of Managed Objects for the LDP Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths</i>	Unsupported tables and objects: <ul style="list-style-type: none"> • mplsMldpInterfaceStatsTable <p>Also, the following fields of the mplsMldpFecUpstreamSessTable are not implemented because these statistics are not currently supported in LDP or PFE:</p> <ul style="list-style-type: none"> • mplsMldpFecUpstreamSessPackets • mplsMldpFecUpstreamSessBytes • mplsMldpFecUpstreamSessDiscontinuityTime 	M Series, MX Series, PTX Series, and T Series
Internet draft draft-ietf-mppls-p2mp-te-mib-09.txt, <i>P2MP MPLS-TE MIB</i> (read-only access)	Unsupported table: <ul style="list-style-type: none"> • mplsTeP2mpTunnelBranchPerfTable 	ACX Series, M Series, MX Series, PTX Series, and T Series
Internet draft draft-ietf-ospf-ospfv3-mib-11.txt, <i>Management Information Base for OSPFv3</i>	Support for ospfv3NbrTable only.	M Series, MX Series, PTX Series, SRX Series, and T Series

Table 7: Standard MIBs supported by Junos OS (*continued*)

Standard MIB	Exceptions	Platforms
Internet draft draft-ietf-ppvpn-mpls-vpn-mib-04.txt, <i>MPLS/BGP Virtual Private Network Management Information Base Using SMIPv2</i>	Supported tables and objects: <ul style="list-style-type: none"> • <code>mplsVpnScalars</code> • <code>mplsVpnVrfTable</code> • <code>mplsVpnPerTable</code> • <code>mplsVpnVrfRouteTargetTable</code> 	M Series, MX Series, PTX Series, and T Series
Internet draft draft-reeder-snmpv3-usm-3desede-00.txt, <i>Extension to the User-Based Security Model (USM) to Support Triple-DES EDE in 'Outside' CBC Mode</i>	No exceptions	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, and T Series

For information about standard SNMP MIB objects, see the [SNMP MIB Explorer](#).

Related Documentation

- [Enterprise-Specific SNMP MIBs Supported by Junos OS on page 21](#)
- [Network Management Administration Guide](#)

Standard SNMP Traps Supported by Junos OS

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

This topic provides the list of standard SNMPv1 and SNMPv2 traps supported by devices running Junos OS. For more information about traps see [SNMP MIB Explorer](#).

Standard SNMP Version 1 Traps

[Table 8 on page 50](#) provides an overview of the standard traps for SNMPv1. The traps are organized first by trap category and then by trap name, and include their enterprise ID, generic trap number, and specific trap number. The system logging severity levels are listed for those traps that have them with their corresponding system log tag. Traps that do not have corresponding system logging severity levels are marked with an en dash (–) in the table.

For more information about system log messages, see the [System Log Explorer](#). For more information about configuring system logging, see the *Junos OS System Basics Configuration Guide*.

Table 8: Standard Supported SNMP Version 1 Traps

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag	Supported On
------------	-----------	---------------	---------------------	----------------------	-------------------------------	------------	--------------

Startup Notifications

Table 8: Standard Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag	Supported On
RFC 1215, <i>Conventions for Defining Traps for Use with the SNMP</i>	authenticationFailure	1.3.6.1.4.1.2636	4	0	Notice	SNMPD_TRAP_GEN_FAILURE	All devices running Junos OS.
	coldStart	1.3.6.1.4.1.2636	0	0	Critical	SNMPD_TRAP_COLD_START	All devices running Junos OS.
	warmStart	1.3.6.1.4.1.2636	1	0	Error	SNMPD_TRAP_WARM_START	All devices running Junos OS.
Link Notifications							
RFC 1215, <i>Conventions for Defining Traps for Use with the SNMP</i>	linkDown	1.3.6.1.4.1.2636	2	0	Warning	SNMP_TRAP_LINK_DOWN	All devices running Junos OS.
	linkUp	1.3.6.1.4.1.2636	3	0	Info	SNMP_TRAP_LINK_UP	All devices running Junos OS.
Remote Operations Notifications							
RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	pingProbeFailed	1.3.6.1.2.1.80.0	6	1	Info	SNMP_TRAP_PING_PROBE_FAILED	All devices running Junos OS.
	pingTestFailed	1.3.6.1.2.1.80.0	6	2	Info	SNMP_TRAP_PING_TEST_FAILED	All devices running Junos OS.
	pingTestCompleted	1.3.6.1.2.1.80.0	6	3	Info	SNMP_TRAP_PING_TEST_COMPLETED	All devices running Junos OS.
	traceRoutePathChange	1.3.6.1.2.1.81.0	6	1	Info	SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE	All devices running Junos OS.
	traceRouteTestFailed	1.3.6.1.2.1.81.0	6	2	Info	SNMP_TRAP_TRACE_ROUTE_TEST_FAILED	All devices running Junos OS.
	traceRouteTestCompleted	1.3.6.1.2.1.81.0	6	3	Info	SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED	All devices running Junos OS.
RMON Alarms							
RFC 2819a, <i>RMON MIB</i>	fallingAlarm	1.3.6.1.2.1.16	6	2	—	—	All devices running Junos OS.
	risingAlarm	1.3.6.1.2.1.16	6	1	—	—	All devices running Junos OS.

Table 8: Standard Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag	Supported On
Routing Notifications							
<i>BGP 4 MIB</i>	bgpEstablished	1.3.6.1.2.1.15.7	6	1	–	–	M, T, MX, J, EX, and SRX Series devices.
	bgpBackwardTransition	1.3.6.1.2.1.15.7	6	2	–	–	M, T, MX, J, EX, and SRX Series devices.
<i>OSPF TRAP MIB</i>	ospfVirtIfStateChange	1.3.6.1.2.1.14.16.2	6	1	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfNbrStateChange	1.3.6.1.2.1.14.16.2	6	2	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfVirtNbrStateChange	1.3.6.1.2.1.14.16.2	6	3	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfIfConfigError	1.3.6.1.2.1.14.16.2	6	4	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfVirtIfConfigError	1.3.6.1.2.1.14.16.2	6	5	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfIfAuthFailure	1.3.6.1.2.1.14.16.2	6	6	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfVirtIfAuthFailure	1.3.6.1.2.1.14.16.2	6	7	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfIfRxBadPacket	1.3.6.1.2.1.14.16.2	6	8	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfVirtIfRxBadPacket	1.3.6.1.2.1.14.16.2	6	9	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfTxRetransmit	1.3.6.1.2.1.14.16.2	6	10	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfVirtIfTxRetransmit	1.3.6.1.2.1.14.16.2	6	11	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfMaxAgeLsa	1.3.6.1.2.1.14.16.2	6	13	–	–	M, T, MX, J, EX, and SRX Series devices.
	ospfIfStateChange	1.3.6.1.2.1.14.16.2	6	16	–	–	M, T, MX, J, EX, and SRX Series devices.

Table 8: Standard Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	Syslog Tag	Supported On
VRRP Notifications							
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	vrrpTrapNewMaster	1.3.6.1.2.1.68	6	1	Warning	VRRPD_NEW_MASTER_TRAP	All devices running Junos OS.
	vrrpTrapAuthFailure	1.3.6.1.2.1.68	6	2	Warning	VRRPD_AUTH_FAILURE_TRAP	All devices running Junos OS.
RFC 6527, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol Version 3 (VRRPv3)</i>	vrrpv3NewMaster	1.3.6.1.2.1.207	6	1	Warning	VRRPD_NEW_MASTER	M and MX
	vrrpv3ProtoError	1.3.6.1.2.1.207	6	2	Warning	VRRPD_V3_PROTO_ERROR	M and MX

Standard SNMP Version 2 Traps

[Table 9 on page 53](#) provides an overview of the standard SNMPv2 traps supported by the Junos OS. The traps are organized first by trap category and then by trap name and include their **snmpTrapOID**. The system logging severity levels are listed for those traps that have them with their corresponding system log tag. Traps that do not have corresponding system logging severity levels are marked with an en dash (–) in the table.

For more information about system log messages, see *System Log Messages Configuration Guide*.

Table 9: Standard Supported SNMP Version 2 Traps

Defined in	Trap Name	snmpTrapOID	System Logging Severity Level	Syslog Tag	Supported On
------------	-----------	-------------	-------------------------------	------------	--------------

Startup Notifications

Table 9: Standard Supported SNMP Version 2 Traps (*continued*)

Defined in	Trap Name	snmpTrapOID	System Logging Severity Level	Syslog Tag	Supported On
RFC 1907, <i>Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)</i>	coldStart	1.3.6.1.6.3.1.1.5.1	Critical	SNMPD_TRAP_COLD_START	All devices running Junos OS.
	warmStart	1.3.6.1.6.3.1.1.5.2	Error	SNMPD_TRAP_WARM_START	All devices running Junos OS.
	authenticationFailure	1.3.6.1.6.3.1.1.5.5	Notice	SNMPD_TRAP_GEN_FAILURE	All devices running Junos OS.
Link Notifications					
RFC 2863, <i>The Interfaces Group MIB</i>	linkDown	1.3.6.1.6.3.1.1.5.3	Warning	SNMP_TRAP_LINK_DOWN	All devices running Junos OS.
	linkUp	1.3.6.1.6.3.1.1.5.4	Info	SNMP_TRAP_LINK_UP	All devices running Junos OS.
Remote Operations Notifications					
RFC 2925, <i>Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations</i>	pingProbeFailed	1.3.6.1.2.1.80.0.1	Info	SNMP_TRAP_PING_PROBE_FAILED	All devices running Junos OS.
	pingTestFailed	1.3.6.1.2.1.80.0.2	Info	SNMP_TRAP_PING_TEST_FAILED	All devices running Junos OS.
	pingTestCompleted	1.3.6.1.2.1.80.0.3	Info	SNMP_TRAP_PING_TEST_COMPLETED	All devices running Junos OS.
	traceRoutePathChange	1.3.6.1.2.1.81.0.1	Info	SNMP_TRAP_TRACE_ROUTE_PATH_CHANGE	All devices running Junos OS.
	traceRouteTestFailed	1.3.6.1.2.1.81.0.2	Info	SNMP_TRAP_TRACE_ROUTE_TEST_FAILED	All devices running Junos OS.
	traceRouteTestCompleted	1.3.6.1.2.1.81.0.3	Info	SNMP_TRAP_TRACE_ROUTE_TEST_COMPLETED	All devices running Junos OS.
RMON Alarms					
RFC 2819a, <i>RMON MIB</i>	fallingAlarm	1.3.6.1.2.1.16.0.1	–	–	All devices running Junos OS.
	risingAlarm	1.3.6.1.2.1.16.0.2	–	–	All devices running Junos OS.

Table 9: Standard Supported SNMP Version 2 Traps (*continued*)

Defined in	Trap Name	snmpTrapOID	System Logging Severity Level	Syslog Tag	Supported On
Routing Notifications					
<i>BGP 4 MIB</i>	bgpEstablished	1.3.6.1.2.1.15.7.1	–	–	All devices running Junos OS.
	bgpBackwardTransition	1.3.6.1.2.1.15.7.2	–	–	All devices running Junos OS.
<i>OSPF Trap MIB</i>	ospfVirtIfStateChange	1.3.6.1.2.1.14.16.2.1	–	–	All devices running Junos OS.
	ospfNbrStateChange	1.3.6.1.2.1.14.16.2.2	–	–	All devices running Junos OS.
	ospfVirtNbrStateChange	1.3.6.1.2.1.14.16.2.3	–	–	All devices running Junos OS.
	ospfIfConfigError	1.3.6.1.2.1.14.16.2.4	–	–	All devices running Junos OS.
	ospfVirtIfConfigError	1.3.6.1.2.1.14.16.2.5	–	–	All devices running Junos OS.
	ospfIfAuthFailure	1.3.6.1.2.1.14.16.2.6	–	–	All devices running Junos OS.
	ospfVirtIfAuthFailure	1.3.6.1.2.1.14.16.2.7	–	–	All devices running Junos OS.
	ospfIfRxBadPacket	1.3.6.1.2.1.14.16.2.8	–	–	All devices running Junos OS.
	ospfVirtIfRxBadPacket	1.3.6.1.2.1.14.16.2.9	–	–	All devices running Junos OS.
	ospfTxRetransmit	1.3.6.1.2.1.14.16.2.10	–	–	All devices running Junos OS.
	ospfVirtIfTxRetransmit	1.3.6.1.2.1.14.16.2.11	–	–	All devices running Junos OS.
	ospfMaxAgeLsa	1.3.6.1.2.1.14.16.2.13	–	–	All devices running Junos OS.
	ospfIfStateChange	1.3.6.1.2.1.14.16.2.16	–	–	All devices running Junos OS.

Table 9: Standard Supported SNMP Version 2 Traps (*continued*)

Defined in	Trap Name	snmpTrapOID	System Logging Severity Level	Syslog Tag	Supported On
MPLS Notifications					
RFC 3812, <i>Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base</i>	mplsTunnelUp				
	mplsTunnelDown				
	mplsTunnelRerouted				
	mplsTunnelReoptimized				
Entity State MIB Notifications					
RFC 4268, <i>Entity State MIB</i>	entStateOperEnabled	1.3.6.1.2.1.131.0.1	Notice	CHASSISD_SNMP_TRAP3	MX240, MX480, and MX960
	entStateOperDisabled	1.3.6.1.2.1.131.0.2	Notice	CHASSISD_SNMP_TRAP3	MX240, MX480, and MX960
L3VPN Notifications					
RFC 4382, <i>MPLS/BGP Layer 3 Virtual Private Network (VPN)</i>	mplsL3VpnVrfUp				
	mplsL3VpnVrfDown				
	mplsL3VpnVrfRouteMidThreshExceeded				
	mplsL3VpnVrfNumVrfRouteMaxThreshExceeded				
	mplsL3VpnNumVrfRouteMaxThreshCleared				
VRRP Notifications					
RFC 2787, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol</i>	vrrpTrapNewMaster	1.3.6.1.2.1.68.0.1	Warning	VRRPD_NEWMASTER_TRAP	All devices running Junos OS.
	vrrpTrapAuthFailure	1.3.6.1.2.1.68.0.2	Warning	VRRPD_AUTH_FAILURE_TRAP	All devices running Junos OS.

Table 9: Standard Supported SNMP Version 2 Traps (*continued*)

Defined in	Trap Name	snmpTrapOID	System Logging Severity Level	Syslog Tag	Supported On
RFC 6527, <i>Definitions of Managed Objects for the Virtual Router Redundancy Protocol Version 3 (VRRPv3)</i>	vrrpv3NewMaster	1.3.6.1.2.1.207.0.1	Warning	VRRPD_NEW_MASTER	M and MX
	vrrpv3ProtoError	1.3.6.1.2.1.207.0.2	Warning	VRRPD_V3_PROTO_ERROR	M and MX

Related Documentation

- [Enterprise-Specific SNMP Traps Supported by Junos OS on page 57](#)
- [Enterprise-Specific SNMP MIBs Supported by Junos OS on page 21](#)
- [Standard SNMP MIBs Supported by Junos OS on page 32](#)
- [Configuring SNMP Trap Options and Groups on a Device Running Junos OS on page 102](#)
- [Managing Traps and Informs](#)

Enterprise-Specific SNMP Traps Supported by Junos OS

Supported Platforms [ACX Series, EX Series, M Series, MX Series, PTX Series, T Series](#)

This topic provides the list of Juniper Networks enterprise-specific SNMPv1 and SNMPv2 traps supported on devices running Junos OS. For more information about traps see [SNMP MIB Explorer](#).

- [Juniper Networks Enterprise-Specific SNMP Version 1 Traps on page 57](#)
- [Juniper Networks Enterprise-Specific SNMP Version 2 Traps on page 63](#)

Juniper Networks Enterprise-Specific SNMP Version 1 Traps

The Junos OS supports enterprise-specific SNMP version 1 traps shown in [Table 10 on page 58](#). The traps are organized first by trap category and then by trap name. The system logging severity levels are listed for those traps that have them. Traps that do not have corresponding system logging severity levels are marked with an en dash (–).

For more information about system log messages, see the *Junos OS System Log Reference for Security Devices*.

Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag	Supported On
Chassis Notifications (Alarm Conditions)							
<i>Chassis MIB</i> (jnx-chassis.mib)	jnxPowerSupplyFailure	1.3.6.1.4.1.2636.4.1	6	1	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFanFailure	1.3.6.1.4.1.2636.4.1	6	2	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxOverTemperature	1.3.6.1.4.1.2636.4.1	6	3	Alert	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxRedundancySwitchOver	1.3.6.1.4.1.2636.4.1	6	4	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruRemoval	1.3.6.1.4.1.2636.4.1	6	5	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruInsertion	1.3.6.1.4.1.2636.4.1	6	6	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruPowerOff	1.3.6.1.4.1.2636.4.1	6	7	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruPowerOn	1.3.6.1.4.1.2636.4.1	6	8	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruFailed	1.3.6.1.4.1.2636.4.1	6	9	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruOffline	1.3.6.1.4.1.2636.4.1	6	10	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruOnline	1.3.6.1.4.1.2636.4.1	6	11	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruCheck	1.3.6.1.4.1.2636.4.1	6	12	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.

Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag	Supported On
	jnxFEBSwitchover	1.3.6.1.4.1.2636.4.1	6	13	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxHardDiskFailed	1.3.6.1.4.1.2636.4.1	6	14	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxHardDiskMissing	1.3.6.1.4.1.2636.4.1	6	15	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxPowerSupplyOk	1.3.6.1.4.1.2636.4.2	6	1	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFanOK	1.3.6.1.4.1.2636.4.2	6	2	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxTemperatureOK	1.3.6.1.4.1.2636.4.2	6	3	Alert	CHASSISD_SNMP_TRAP	All devices running Junos OS.
Configuration Notifications							
<i>Configuration Management MIB (jnx-configmgmt.mib)</i>	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5	6	1	—	—	All devices running Junos OS.
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5	6	2	—	—	All devices running Junos OS.

Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag	Supported On
<i>Flow Collection Services MIB (jnx-coll.mib)</i>	jnxCollUnavailableDest	1.3.6.1.4.1.2636.4.8	6	1	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollUnavailableDestCleared	1.3.6.1.4.1.2636.4.8	6	2	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollUnsuccessfulTransfer	1.3.6.1.4.1.2636.4.8	6	3	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollFlowOverload	1.3.6.1.4.1.2636.4.8	6	4	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollFlowOverloadCleared	1.3.6.1.4.1.2636.4.8	6	5	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollMemoryUnavailable	1.3.6.1.4.1.2636.4.8	6	6	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollMemoryAvailable	1.3.6.1.4.1.2636.4.8	6	7	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollFtpSwitchover	1.3.6.1.4.1.2636.4.8	6	8	–	–	Devices that run Junos OS and have collector PICs installed.

Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag	Supported On
<i>Passive Monitoring MIB</i> (jnx-pmonmib)	jnxPMonOverloadSet	1.3.6.1.4.1.2636.4.7.0.1	6	1	–	–	Devices that run Junos OS and have PICs that support passive monitoring installed.
	jnxPMonOverloadCleared	1.3.6.1.4.1.2636.4.7.0.2	6	2	–	–	Devices that run Junos OS and have PICs that support passive monitoring installed.
<i>SONET APS MIB</i> (jnx-sonetaps.mib)	apsEventChannelMismatch	1.3.6.1.4.1.2636.3.24.2	6	3	–	–	Devices that run Junos OS and have SONET PICs installed.
	apsEventPSBF	1.3.6.1.4.1.2636.3.24.2	6	4	–	–	Devices that run Junos OS and have SONET PICs installed.
	apsEventFEPLF	1.3.6.1.4.1.2636.3.24.2	6	5	–	–	Devices that run Junos OS and have SONET PICs installed.
Remote Operations							
<i>PING MIB</i> (jnx-ping.mib)	jnxPingRttThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	1	–	–	All devices running Junos OS.
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	2	–	–	All devices running Junos OS.
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	3	–	–	All devices running Junos OS.
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	4	–	–	All devices running Junos OS.

Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag	Supported On
	jnxPingEgressStdDevThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	5	–	–	All devices running Junos OS.
	jnxPingEgressJitterThresholdExceeded	1.3.6.1.4.1.2636.4.9	6	6	–	–	All devices running Junos OS.
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	7	–	–	All devices running Junos OS.
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	8	–	–	All devices running Junos OS.
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9	6	9	–	–	All devices running Junos OS.
Routing Notifications							
<i>BFD Experimental MIB (jnx-bfd-exp.mib)</i>	bfdSessUp	1.3.6.1.4.1.2636.5.3.1	6	1	–	–	All devices running Junos OS.
	bfdSessDown	1.3.6.1.4.1.2636.5.3.1	6	2	–	–	All devices running Junos OS.
<i>LDP MIB (jnx-ldp.mib)</i>	jnxLdpLspUp	1.3.6.1.4.1.2636.4.4	6	1	–	–	M, T, and MX Series routers.
	jnxLdpLspDown	1.3.6.1.4.1.2636.4.4	6	2	–	–	M, T, and MX Series routers.
	jnxLdpSesUp	1.3.6.1.4.1.2636.4.4	6	3	–	–	M, T, and MX Series routers.
	jnxLdpSesDown	1.3.6.1.4.1.2636.4.4	6	4	–	–	M, T, and MX Series routers.

Table 10: Juniper Networks Enterprise-Specific Supported SNMP Version 1 Traps (*continued*)

Defined in	Trap Name	Enterprise ID	Generic Trap Number	Specific Trap Number	System Logging Severity Level	System Log Tag	Supported On
<i>MPLS MIB</i> (<i>jnx-mpls.mib</i>)	mplsLspUp (Deprecated)	1.3.6.1.4.1.2636.3.2.4	6	1	—	—	
	mplsLspDown (Deprecated)	1.3.6.1.4.1.2636.3.2.4	6	2	—	—	
	mplsLspChange (Deprecated)	1.3.6.1.4.1.2636.3.2.4	6	3	—	—	
	mplsLspPathDown (Deprecated)	1.3.6.1.4.1.2636.3.2.4	6	4	—	—	
<i>VPN MIB</i> (<i>jnx-vpn.mib</i>)	jnxVpnIfUp	1.3.6.1.4.1.2636.3.2.6	6	1	—	—	M, T, and MX Series routers.
	jnxVpnIfDown	1.3.6.1.4.1.2636.3.2.6	6	2	—	—	M, T, and MX Series routers.
	jnxVpnPwUp	1.3.6.1.4.1.2636.3.2.6	6	3	—	—	M, T, and MX Series routers.
	jnxVpnPwDown	1.3.6.1.4.1.2636.3.2.6	6	4	—	—	M, T, and MX Series routers.
RMON Alarms							
<i>RMON MIB</i> (<i>jnx-rmon.mib</i>)	jnxRmonAlarmGetFailure	1.3.6.1.4.1.2636.4.3	6	1	—	—	All devices running Junos OS.
	jnxRmonGetOk	1.3.6.1.4.1.2636.4.3	6	2	—	—	All devices running Junos OS.
SONET Alarms							
<i>SONET MIB</i> (<i>jnx-sonet.mib</i>)	jnxSonetAlarmSet	1.3.6.1.4.1.2636.4.6	6	1	—	—	Devices that run Junos OS and have SONET PICs installed.
	jnxSonetAlarmCleared	1.3.6.1.4.1.2636.4.6	6	2	—	—	Devices that run Junos OS and have SONET PICs installed.

Juniper Networks Enterprise-Specific SNMP Version 2 Traps

The Junos OS supports the enterprise-specific SNMP version 2 traps shown in [Table 11 on page 64](#). The traps are organized first by trap category and then by trap name.

The system logging severity levels are listed for those traps that have them. Traps that do not have corresponding system logging severity levels are marked with an en dash (–).

For more information about system messages, see the [System Log Explorer](#). For more information about configuring system logging, see the [Junos OS Administration Library for Routing Devices](#).

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
Chassis (Alarm Conditions) Notifications					
<i>Chassis MIB</i> (jnx-chassis.mib)	jnxPowerSupplyFailure	1.3.6.1.4.1.2636.4.1.1	Alert	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFanFailure	1.3.6.1.4.1.2636.4.1.2	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxOverTemperature	1.3.6.1.4.1.2636.4.1.3	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruNotifAdminStatus		Notice		
	jnxFruNotifMismatch		Notice		
	jnxFruNotifOperStatus		Notice		
	jnxRedundancySwitchOver	1.3.6.1.4.1.2636.4.1.4	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruRemoval	1.3.6.1.4.1.2636.4.1.5	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruInsertion	1.3.6.1.4.1.2636.4.1.6	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruPowerOff	1.3.6.1.4.1.2636.4.1.7	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruPowerOn	1.3.6.1.4.1.2636.4.1.8	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruFailed	1.3.6.1.4.1.2636.4.1.9	Warning	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruOffline	1.3.6.1.4.1.2636.4.1.10	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
	jnxFruOnline	1.3.6.1.4.1.2636.4.1.11	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFruCheck	1.3.6.1.4.1.2636.4.1.12	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxFEBSwitchover	1.3.6.1.4.1.2636.4.1.13	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxHardDiskFailed	1.3.6.1.4.1.2636.4.1.14	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
	jnxHardDiskMissing	1.3.6.1.4.1.2636.4.1.15	Notice	CHASSISD_SNMP_TRAP	All devices running Junos OS.
jnxPowerSupplyOK	1.3.6.1.4.1.2636.4.2.1	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.	
jnxFanOK	1.3.6.1.4.1.2636.4.2.2	Critical	CHASSISD_SNMP_TRAP	All devices running Junos OS.	
jnxTemperatureOK	1.3.6.1.4.1.2636.4.2.3	Alert	CHASSISD_SNMP_TRAP	All devices running Junos OS.	
Configuration Notifications					
<i>Configuration Management MIB (jnx-cfgmgmt.mib)</i>	jnxCmCfgChange	1.3.6.1.4.1.2636.4.5.0.1	–	–	All devices running Junos OS.
	jnxCmRescueChange	1.3.6.1.4.1.2636.4.5.0.2	–	–	All devices running Junos OS.
Link Notifications					

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
<i>Flow Collection Services MIB (jnx-coll.mib)</i>	jnxCollUnavailableDest	1.3.6.1.4.1.2636.4.8.0.1	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollUnavailableDestCleared	1.3.6.1.4.1.2636.4.8.0.2	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollUnsuccessfulTransfer	1.3.6.1.4.1.2636.4.8.0.3	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollFlowOverload	1.3.6.1.4.1.2636.4.8.0.4	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollFlowOverloadCleared	1.3.6.1.4.1.2636.4.8.0.5	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollMemoryUnavailable	1.3.6.1.4.1.2636.4.8.0.6	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollMemoryAvailable	1.3.6.1.4.1.2636.4.8.0.7	–	–	Devices that run Junos OS and have collector PICs installed.
	jnxCollFtpSwitchover	1.3.6.1.4.1.2636.4.8.0.8	–	–	Devices that run Junos OS and have collector PICs installed.
<i>PMON MIB (jnx-pmon.mib)</i>	jnxPMonOverloadSet	1.3.6.1.4.1.2636.4.7.0.1	–	–	Devices that run Junos OS and have PICs that support passive monitoring installed.
	jnxPMonOverloadCleared	1.3.6.1.4.1.2636.4.7.0.2	–	–	Devices that run Junos OS and have PICs that support passive monitoring installed.

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
<i>SONET APS MIB (jnx-sonetaps.mib)</i>	apsEventChannelMismatch	1.3.6.1.4.1.2636.3.24.2.0.3	–	–	Devices that run Junos OS and have SONET PICs installed.
	apsEventPSBF	1.3.6.1.4.1.2636.3.24.2.0.4	–	–	Devices that run Junos OS and have SONET PICs installed.
	apsEventFEPLF	1.3.6.1.4.1.2636.3.24.2.0.5	–	–	Devices that run Junos OS and have SONET PICs installed.
Remote Operations Notifications					
<i>PING MIB (jnx-ping.mib)</i>	jnxPingRttThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.1	–	–	All devices running Junos OS.
	jnxPingRttStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.2	–	–	All devices running Junos OS.
	jnxPingRttJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.3	–	–	All devices running Junos OS.
	jnxPingEgressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.4	–	–	All devices running Junos OS.
	jnxPingEgressStdDevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.5	–	–	All devices running Junos OS.
	jnxPingEgressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.6	–	–	All devices running Junos OS.
	jnxPingIngressThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.7	–	–	All devices running Junos OS.
	jnxPingIngressStddevThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.8	–	–	All devices running Junos OS.
	jnxPingIngressJitterThreshold Exceeded	1.3.6.1.4.1.2636.4.9.0.9	–	–	All devices running Junos OS.
Routing Notifications					
<i>BFD Experimental MIB (jnx-bfd-exp.mib)</i>	bfdSessUp	1.3.6.1.4.1.2636.5.3.1.0.1	–	–	All devices running Junos OS.
	bfdSessDown	1.3.6.1.4.1.2636.5.3.1.0.2	–	–	All devices running Junos OS.

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
<i>BGP4 V2 MIB</i> (jnx-bgpmib2.mib)	jnxBgpM2Established	1.3.6.1.4.1.2636.5.1.1.0.1	–	–	All devices running Junos OS.
	jnxBgpM2BackwardTransition	1.3.6.1.4.1.2636.5.1.1.0.2	–	–	All devices running Junos OS.
<i>DHCP MIB</i> (jnx-dhcp.mib)	jnxJdhcpLocalServer DuplicateClient	1.3.6.1.4.1.2636.3.6.1.6.1.3.1	–	–	All devices running Junos OS.
	jnxJdhcpLocalServer InterfaceLimitExceeded	1.3.6.1.4.1.2636.3.6.1.6.1.3.2	–	–	All devices running Junos OS.
	jnxJdhcpLocalServer InterfaceLimitAbated	1.3.6.1.4.1.2636.3.6.1.6.1.3.3	–	–	All devices running Junos OS.
	jnxJdhcpLocalServer Health	1.3.6.1.4.1.2636.3.6.1.6.1.3.4	–	–	All devices running Junos OS.
	jnxJdhcpRelayInterface LimitExceeded	1.3.6.1.4.1.2636.3.6.1.6.1.2.3.1	–	–	All devices running Junos OS.
	jnxJdhcpRelayInterface LimitAbated	1.3.6.1.4.1.2636.3.6.1.6.1.2.3.2	–	–	All devices running Junos OS.
<i>DHCPv6 MIB</i> (jnx-dhcpv6.mib)	jnxJdhcpv6LocalServer InterfaceLimitExceeded	1.3.6.1.4.1.2636.3.6.2.6.2.2.3.1	–	–	All devices running Junos OS.
	jnxJdhcpv6LocalServer InterfaceLimitAbated	1.3.6.1.4.1.2636.3.6.2.6.2.2.3.2	–	–	All devices running Junos OS.
	jnxJdhcpv6LocalServer Health	1.3.6.1.4.1.2636.3.6.2.6.2.2.3.3	–	–	All devices running Junos OS.
<i>LDP MIB</i> (jnx-ldp.mib)	jnxLdpLspUp	1.3.6.1.4.1.2636.4.4.0.1	–	–	M, T, and MX Series routers.
	jnxLdpLspDown	1.3.6.1.4.1.2636.4.4.0.2	–	–	M, T, and MX Series routers.
	jnxLdpSesUp	1.3.6.1.4.1.2636.4.4.0.3	–	–	M, T, and MX Series routers.
	jnxLdpSesDown	1.3.6.1.4.1.2636.4.4.0.4	–	–	M, T, and MX Series routers.

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
MPLS MIB (jnx-mpls.mib)	mplsLspUp (Deprecated)	1.3.6.1.4.1.2636.3.2.4.1	–	–	
	mplsLspInfoUp	1.3.6.1.4.1.2636.3.2.0.1	–	–	M, T, and MX Series routers.
	mplsLspDown (Deprecated)	1.3.6.1.4.1.2636.3.2.4.2	–	–	
	mplsLspInfoDown	1.3.6.1.4.1.2636.3.2.0.2	–	–	M, T, and MX Series routers.
	mplsLspChange (Deprecated)	1.3.6.1.4.1.2636.3.2.4.3	–	–	
	mplsLspInfoChange	1.3.6.1.4.1.2636.3.2.0.3	–	–	M, T, and MX Series routers.
	mplsLspPathDown (Deprecated)	1.3.6.1.4.1.2636.3.2.4.4	–	–	
	mplsLspInfoPathDown	1.3.6.1.4.1.2636.3.2.0.4	–	–	M, T, and MX Series routers.
mplsLspInfoPathUp	1.3.6.1.4.1.2636.3.2.0.5	–	–	M, T, and MX Series routers.	
VPN MIB (jnx-vpn.mib)	jnxVpnIfUp	1.3.6.1.4.1.2636.3.26.0.1	–	–	M, T, and MX Series routers.
	jnxVpnIfDown	1.3.6.1.4.1.2636.3.26.0.2	–	–	M, T, and MX Series routers.
	jnxVpnPwUp	1.3.6.1.4.1.2636.3.26.0.3	–	–	M, T, and MX Series routers.
	jnxVpnPwDown	1.3.6.1.4.1.2636.3.26.0.4	–	–	M, T, and MX Series routers.

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
AAA MIB (jnx-user-aaa.mib)	jnxAccessAuthAddressPoolHighThreshold	1.3.6.1.4.1.2636.3.51.1.0.5	–	–	SRX Series devices.
	jnxAccessAuthAddressPoolAbateThreshold	1.3.6.1.4.1.2636.3.51.1.0.6	–	–	SRX Series devices.
	jnxAccessAuthAddressPoolOutOfAddresses	1.3.6.1.4.1.2636.3.51.1.0.7	–	–	SRX Series devices.
	jnxAccessAuthAddressPoolOutOfMemory	1.3.6.1.4.1.2636.3.51.1.0.8	–	–	SRX Series devices.
	jnxAccessAuthService Up	1.3.6.1.4.1.2636.3.51.1.0.1	–	–	SRX Series devices.
	jnxAccessAuthService Down	1.3.6.1.4.1.2636.3.51.1.0.2	–	–	SRX Series devices.
	jnxAccessAuthServer Disabled	1.3.6.1.4.1.2636.3.51.1.0.3	–	–	SRX Series devices.
	jnxAccessAuthServer Enabled	1.3.6.1.4.1.2636.3.51.1.0.4	–	–	SRX Series devices.
Access Authentication Methods MIB (jnx-js-auth.mib)	jnxJsFwAuthFailure	1.3.6.1.4.1.2636.3.39.1.2.1.0.1	–	–	SRX Series devices.
	jnxJsFwAuthServiceUp	1.3.6.1.4.1.2636.3.39.1.2.1.0.2	–	–	SRX Series devices.
	jnxJsFwAuthServiceDown	1.3.6.1.4.1.2636.3.39.1.2.1.0.3	–	–	SRX Series devices.
	jnxJsFwAuthCapacityExceeded	1.3.6.1.4.1.2636.3.39.1.2.1.0.4	–	–	SRX Series devices.
Network Address Translation Resources Monitoring MIB (jnxNatMIB)	jnxJsNatAddrPoolThresholdStatus	1.3.6.1.4.1.2636.3.39.1.7.1.0.1	–	–	SRX Series devices.
	jnxNatAddrPoolUtil	1.3.6.1.4.1.2636.3.59.1.2.1	–	–	M Series and MX Series routers
	jnxNatTrapSrcPoolName	1.3.6.1.4.1.2636.3.59.1.2.2	–	–	M Series and MX Series routers
	jnxNatAddrPoolThresholdStatus	1.3.6.1.4.1.2636.3.59.1.0.1	–	–	M Series and MX Series routers

Table 11: Juniper Networks Enterprise-Specific Supported SNMP Version 2 Traps (*continued*)

Source MIB	Trap Name	snmpTrapOID	System Logging Severity Level	System Log Tag	Supported On
<i>Network Address Translation MIB</i> (jnx-js-nat.mib)	jnxJsScreen Attack	1.3.6.1.4.1.2636.3.39.1.8.1.0.1	Warning	RT_SCREEN_ICMP, RT_SCREEN_IP, RT_SCREEN_SESSION_LIMIT, RT_SCREEN_TCP, RT_SCREEN_UDP	SRX Series devices.
<i>Security Screening Objects MIB</i> (jnx-js-screening.mib)	jnxJsScreenCfg Change	1.3.6.1.4.1.2636.3.39.1.8.1.0.2	–	–	SRX Series devices.
RMON Alarms					
<i>RMON MIB</i> (jnx-rmon.mib)	jnxRmonGetOk	1.3.6.1.4.1.2636.4.3.0.2	–	–	All devices running Junos OS.
SONET Alarms					
<i>SONET MIB</i> (jnx-sonet.mib)	jnxSonetAlarm Cleared	1.3.6.1.4.1.2636.4.6.0.2	–	–	Devices that run Junos OS and have SONET PICs installed.

- Related Documentation**
- [Standard SNMP Traps Supported by Junos OS on page 50](#)
 - [Standard SNMP MIBs Supported by Junos OS on page 32](#)
 - [Enterprise-Specific SNMP MIBs Supported by Junos OS on page 21](#)
 - [Configuring SNMP Trap Options and Groups on a Device Running Junos OS on page 102](#)
 - [Managing Traps and Informs](#)

CHAPTER 5

Loading MIB Files to a Network Management System

- Loading MIB Files to a Network Management System on page 73

Loading MIB Files to a Network Management System

Supported Platforms [ACX Series](#)

For your network management system (NMS) to identify and understand the MIB objects used by the Junos OS, you must first load the MIB files to your NMS using a MIB compiler. A MIB compiler is a utility that parses the MIB information such as the MIB object name, IDs, and data type for the NMS.

You can download the Junos MIB package from the Junos OS Enterprise MIBs index at http://www.juniper.net/techpubs/en_US/release-independent/junos/mibs/mibs.html. The Junos MIB package is available in **.zip** and **.tar** packages. You can download the appropriate format based on your requirements.

The Junos MIB package contains two folders: **StandardMibs** and **JuniperMibs**. The **StandardMibs** folder contains the standard MIBs and RFCs that are supported on devices running the Junos OS, whereas the **JuniperMibs** folder contains the Juniper Networks enterprise-specific MIBs.

To load MIB files that are required for managing and monitoring devices running the Junos OS:

1. Go to the SNMP MIB Explorer Download page for Juniper Networks SNMP MIB packages (<https://apps.juniper.net/mib-explorer/download.jsp#product=Junos%20OS>).
2. Click the **TAR** or **ZIP** link under the appropriate release heading to download the Junos MIB package for that release.
3. Decompress the file (**.tar** or **.zip**) using an appropriate utility.
4. Load the standard MIB files (from the **StandardMibs** folder) in the following order:



NOTE: Some of the MIB compilers that are commonly used have the standard MIBs preloaded on them. If the standard MIBs are already loaded on the MIB compiler that you are using, skip this step and proceed to Step 7.

- a. `mib-SNMPv2-SMI.txt`
 - b. `mib-SNMPv2-TC.txt`
 - c. `mib-IANAIfType-MIB.txt`
 - d. `mib-IANA-RTPROTO-MIB.txt`
 - e. `mib-rfc1907.txt`
 - f. `mib-rfc2011a.txt`
 - g. `mib-rfc2012a.txt`
 - h. `mib-rfc2013a.txt`
 - i. `mib-rfc2863a.txt`
5. Load the remaining standard MIB files.



NOTE: You must follow the order specified in this procedure, and ensure that all standard MIBs are loaded before you load the enterprise-specific MIBs. There might be dependencies that require a particular MIB to be present on the compiler before loading some other MIB. You can find such dependencies listed in the **IMPORT** section of the MIB file.

6. Load the Juniper Networks enterprise-specific SMI MIB, `mib-jnx-smi.txt`, and the following optional SMI MIBs based on your requirements:
- `mib-jnx-js-smi.txt`—(Optional) For Juniper Security MIB tree objects
 - `mib-jnx-ex-smi.txt`—(Optional) For EX Series Ethernet Switches
 - `mib-jnx-exp.txt`—(Recommended) For Juniper Networks experimental MIB objects
7. Load the remaining enterprise-specific MIBs from the **JuniperMibs** folder.



TIP: While loading a MIB file, if the compiler returns an error message saying that any of the objects is undefined, open the MIB file using a text editor and ensure that all the MIB files listed in the **IMPORT** section are loaded on the compiler. If any of the MIB files listed in the **IMPORT** section is not loaded on the compiler, load that MIB file, and then try to load the MIB file that failed to load.

For example, the enterprise-specific PING MIB, `mib-jnx-ping.txt`, has dependencies on RFC 2925, DiSMAN-PING-MIB, `mib-rfc2925a.txt`. If you try to load `mib-jnx-ping.txt` before loading `mib-rfc2925a.txt`, the compiler returns an error message saying that certain objects in `mib-jnx-ping.txt` are undefined. Load `mib-rfc2925a.txt`, and then try to load `mib-jnx-ping.txt`. The enterprise-specific PING MIB, `mib-jnx-ping.txt`, then loads without any issue.

- Related Documentation**
- [Standard SNMP MIBs Supported by Junos OS on page 32](#)
 - [Enterprise-Specific SNMP MIBs Supported by Junos OS on page 21](#)

CHAPTER 6

Configuring SNMP

- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Optimizing the Network Management System Configuration for the Best Results on page 81](#)
- [Configuring Options on Managed Devices for Better SNMP Response Time on page 83](#)
- [Configuring SNMP on Devices Running Junos OS on page 85](#)
- [Configuring the System Contact on a Device Running Junos OS on page 88](#)
- [Configuring the System Location for a Device Running Junos OS on page 89](#)
- [Configuring the System Description on a Device Running Junos OS on page 89](#)
- [Configuring SNMP Details on page 90](#)
- [Configuring a Different System Name on page 91](#)
- [Configuring the Commit Delay Timer on page 92](#)
- [Filtering Duplicate SNMP Requests on page 93](#)
- [Configuring SNMP Communities on page 93](#)
- [Examples: Configuring the SNMP Community String on page 96](#)
- [Adding a Group of Clients to an SNMP Community on page 97](#)
- [Configuring a Proxy SNMP Agent on page 98](#)
- [Configuring SNMP Traps on page 100](#)
- [Configuring SNMP Trap Options and Groups on a Device Running Junos OS on page 102](#)
- [Configuring SNMP Trap Options on page 103](#)
- [Configuring SNMP Trap Groups on page 106](#)
- [Example: Configuring SNMP Trap Groups on page 109](#)
- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 109](#)
- [Example: Configuring Secured Access List Checking on page 110](#)
- [Filtering Interface Information Out of SNMP Get and GetNext Output on page 110](#)
- [Configuring MIB Views on page 111](#)
- [Configuring Ping Proxy MIB on page 113](#)

Configuration Statements at the [edit snmp] Hierarchy Level

Supported Platforms ACX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX

This topic shows all possible configuration statements at the **[edit snmp]** hierarchy level and their level in the configuration hierarchy. When you are configuring Junos OS, your current hierarchy level is shown in the banner on the line preceding the **user@host#** prompt.

```
[edit]
snmp {
  alarm-management {
    alarm-list-name list-name {
      alarm-id id {
        alarm-state state {
          description alarm-description;
          notification-id notification-id-of-alarm;
          resource-prefix alarm-resource-prefix;
          varbind-index varbind-index-in-alarm-varbind-list;
          varbind-subtree alarm-varbind-subtree;
          varbind-value alarm-varbind-value;
        }
      }
    }
  }
  client-list client-list-name {
    ip-addresses;
  }
  community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
      address <restrict>;
    }
    logical-system logical-system-name {
      routing-instance routing-instance-name;
      clients {
        address <restrict>;
      }
    }
    routing-instance routing-instance-name {
      clients {
        address <restrict>;
      }
    }
    view view-name;
  }
  contact contact;
  description description;
  engine-id {
    (local engine-id | use-default-ip-address | use-mac-address);
  }
  filter-duplicates;
  interface [ interface-names ];
```



```

location location;
name name;
nonvolatile {
    commit-delay seconds;
}
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        falling-threshold-interval seconds;
        interval seconds;
        request-type (get-next-request | get-request | walk-request);
        rising-event-index index;
        rising-threshold integer;
        sample-type type;
        startup-alarm alarm;
        syslog-subtag syslog-subtag;
        variable oid-variable;
    }
    event index {
        community community-name;
        description description;
        type type;
    }
}
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
        regular-expression>;
    flag flag;
    memory-trace;
    no-remote-trace;
    no-default-memory-trace;
}
trap-group group-name {
    categories {
        category;
    }
    destination-port port-number;
    routing-instance instance;
    logical-system logical-system-name;
    targets {
        address;
    }
    version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    source-address address;
    enterprise-oid;
    logical-system logical-system-name {
        routing-instance routing-instance-name {
            source-address address;
        }
    }
}
routing-instance routing-instance-name {

```

```
        source-address address;
    }
}
v3 {
    notify name {
        tag tag-name;
        type (trap | inform);
    }
    notify-filter profile-name {
        oid oid (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system logical-system;
        port port-number;
        retry-count number;
        routing-instance instance;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | v3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
}
usm {
    local-engine {
        user username {
            authentication-md5 {
                authentication-password authentication-password;
            }
            authentication-none;
            authentication-sha {
                authentication-password authentication-password;
            }
            privacy-3des {
                privacy-password privacy-password;
            }
            privacy-aes128 {
                privacy-password privacy-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
            privacy-none;
        }
    }
}
```

```

    }
  }
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix){
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
view view-name {
  oid object-identifier (include | exclude);
}
}

```

- Related Documentation**
- [Understanding the SNMP Implementation in Junos OS](#)
 - [Configuring SNMP on a Device Running Junos OS](#)

Optimizing the Network Management System Configuration for the Best Results

Supported Platforms ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX

You can modify your network management system configuration to optimize the response time for SNMP queries. The following sections contain a few tips on how you can configure the network management system:

- [Changing the Polling Method from Column-by-Column to Row-by-Row on page 82](#)
- [Reducing the Number of Variable Bindings per PDU on page 82](#)
- [Increasing Timeout Values in Polling and Discovery Intervals on page 82](#)
- [Reducing Incoming Packet Rate at the snmpd on page 82](#)

Changing the Polling Method from Column-by-Column to Row-by-Row

You can configure the network management system to use the row-by-row method for SNMP data polling. It has been proven that the row-by-row and multiple row-by-multiple-row polling methods are more efficient than column-by-column polling. By configuring the network management system to use the row-by-row data polling method, you can ensure that data for only one interface is polled in a request instead of a single request polling data for multiple interfaces, as is the case with column-by-column polling. Row-by-row polling also reduces the risk of requests timing out.

Reducing the Number of Variable Bindings per PDU

By reducing the number of variable bindings per protocol data unit (PDU), you can improve the response time for SNMP requests. A request that polls for data related to multiple objects, which are mapped to different index entries, translates into multiple requests at the device-end because the subagent might have to poll different modules to obtain data that are linked to different index entries. The recommended method is to ensure that a request has only objects that are linked to one index entry instead of multiple objects linked to different index entries.



NOTE: If responses from a device are slow, avoid using the **GetBulk** option for the device, because a **GetBulk** request might contain objects that are linked to various index entries and might further increase the response time.

Increasing Timeout Values in Polling and Discovery Intervals

By increasing the timeout values for polling and discovery intervals, you can increase the queuing time at the device end and reduce the number of throttle drops that occur because of the request timing out.

Reducing Incoming Packet Rate at the snmpd

By reducing the frequency of sending SNMP requests to a device, you can reduce the risk of SNMP requests piling up at any particular device. Apart from reducing the frequency of sending SNMP requests to a device, you can also increase the polling interval, control the use of **GetNext** requests, and reduce the number of polling stations per device.

Related Documentation

- [Understanding SNMP Implementation in Junos OS on page 15](#)
- [Configuring SNMP on Devices Running Junos OS on page 85](#)
- [Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS on page 195](#)
- [Configuring Options on Managed Devices for Better SNMP Response Time on page 83](#)
- [Managing Traps and Informs](#)
- [Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage](#)

Configuring Options on Managed Devices for Better SNMP Response Time

Supported Platforms ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX

The following sections contain information about configuration options on the managed devices that can enhance SNMP performance:

- [Enabling the stats-cache-lifetime Option on page 83](#)
- [Filtering Out Duplicate SNMP Requests on page 83](#)
- [Excluding Interfaces That Are Slow in Responding to SNMP Queries on page 84](#)

Enabling the stats-cache-lifetime Option

The Junos OS provides you with an option to configure the length of time an SNMP request stays active and queued so as to reduce the possibility of request drops during slow response times. You can use the **stats-cache-lifetime seconds** option at the **[edit snmp]** hierarchy level to specify the length of time that an SNMP request remains queued. The recommended value for the **stats-cache-lifetime** option is in the range of 30 to 60 seconds.



NOTE: The **set snmp stats-cache-lifetime seconds** command is a hidden command and is supported only on devices running Junos OS Release 9.3 and later.

Filtering Out Duplicate SNMP Requests

If a network management station retransmits a **Get**, **GetNext**, or **GetBulk** SNMP request too frequently to a device, that request might interfere with the processing of previous requests and slow down the response time of the agent. Filtering these duplicate requests improves the response time of the SNMP agent. The Junos OS enables you to filter out duplicate **Get**, **GetNext**, and **GetBulk** SNMP requests. The Junos OS uses the following information to determine if an SNMP request is a duplicate:

- Source IP address of the SNMP request
- Source UDP port of the SNMP request
- Request ID of the SNMP request



NOTE: By default, filtering of duplicate SNMP requests is disabled on devices running the Junos OS.

To enable filtering of duplicate SNMP requests on devices running the Junos OS, include the **filter-duplicates** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
filter-duplicates;
```

Excluding Interfaces That Are Slow in Responding to SNMP Queries

An interface that is slow in responding to SNMP requests for interface statistics can delay kernel responses to SNMP requests. You can review the mib2d log file to find out how long the kernel takes to respond to various SNMP requests. For more information about reviewing the log file for kernel response data, see “Checking Kernel and Packet Forwarding Engine Response” under “[Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS](#)” on page 195. If you notice that a particular interface is slow in responding, and think that it is slowing down the kernel from responding to SNMP requests, exclude that interface from the SNMP queries to the device. You can exclude an interface from the SNMP queries either by configuring the **filter-interface** statement or by modifying the SNMP view settings.

The following example shows a sample configuration for excluding interfaces from the SNMP **Get**, **GetNext**, and **Set** operations:

```
[edit]
snmp {
  filter-interfaces {
    interfaces { # exclude the specified interfaces
      interface1;
      interface2;
    }
    all-internal-interfaces; # exclude all internal interfaces
  }
}
```

The following example shows the SNMP view configuration for excluding the interface with an interface index (ifIndex) value of 312 from a request for information related to the ifTable and ifXtable objects:

```
[edit snmp]
view test {
  oid .1 include;
  oid ifTable.1.*:312 exclude;
  oid ifXTable.1.*:312 exclude
}
```

Alternatively, you can take the interface that is slow in responding offline.

Related Documentation

- [Understanding SNMP Implementation in Junos OS on page 15](#)
- [Configuring SNMP on Devices Running Junos OS on page 85](#)
- [Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS on page 195](#)
- [Optimizing the Network Management System Configuration for the Best Results on page 81](#)
- *Managing Traps and Informs*
- *Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage*

Configuring SNMP on Devices Running Junos OS

Supported Platforms ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX

The following sections contain information about basic SNMP configuration and a few examples of configuring the basic SNMP operations on devices running Junos OS:

- [Configuring Basic Settings for SNMPv1 and SNMPv2 on page 85](#)
- [Configuring Basic Settings for SNMPv3 on page 85](#)
- [Configuring System Name, Location, Description, and Contact Information on page 87](#)

Configuring Basic Settings for SNMPv1 and SNMPv2

By default, SNMP is not enabled on devices running Junos OS. To enable SNMP on devices running Junos OS, include the **community public** statement at the **[edit snmp]** hierarchy level.

Enabling SNMPv1 and SNMPv2 Get and GetNext Operations

```
[edit]
snmp {
  community public;
}
```

A community that is defined as public grants access to all MIB data to any client.

To enable SNMPv1 and SNMPv2 **Set** operations on the device, you must include the following statements at the **[edit snmp]** hierarchy level:

Enabling SNMPv1 and SNMPv2 Set Operations

```
[edit snmp]
view all {
  oid .1;
}
community private {
  view all;
  authorization read-write;
}
```

The following example shows the basic minimum configuration for SNMPv1 and SNMPv2 traps on a device:

Configuring SNMPv1 and SNMPv2 Traps

```
[edit snmp]
trap-group jnpr {
  targets {
    192.168.69.179;
  }
}
```

Configuring Basic Settings for SNMPv3

The following example shows the minimum SNMPv3 configuration for enabling **Get**, **GetNext**, and **Set** operations on a device (note that the configuration has authentication set to **md5** and privacy to **none**):

Enabling SNMPv3 Get, GetNext, and Set Operations

```
[edit snmp]
v3 {
  usm {
    local-engine {
      user jnpruser {
        authentication-md5 {
          authentication-key "$9$guaDiQFnAuOQzevMWx7ikqP"; ## SECRET-DATA
        }
        privacy-none;
      }
    }
  }
  vacm {
    security-to-group {
      security-model usm {
        security-name jnpruser {
          group grpnm;
        }
      }
    }
  }
  access {
    group grpnm {
      default-context-prefix {
        security-model any {
          security-level authentication {
            read-view all;
            write-view all;
          }
        }
      }
    }
  }
}
view all {
  oid .1;
}
```

The following example shows the basic configuration for SNMPv3 informs on a device (the configuration has authentication and privacy set to **none**):

Configuring SNMPv3 Informs

```
[edit snmp]
v3 {
  usm {
    remote-engine 00000063200133a2c0a845c3 {
      user RU2_v3_sha_none {
        authentication-none;
        privacy-none;
      }
    }
  }
  vacm {
    security-to-group {
      security-model usm {
        security-name RU2_v3_sha_none {
          group g1_usm_auth;
        }
      }
    }
  }
}
```



```

    }
  }
}
access {
  group g1_usm_auth {
    default-context-prefix {
      security-model usm {
        security-level authentication {
          read-view all;
          write-view all;
          notify-view all;
        }
      }
    }
  }
}
}
}
target-address TA2_v3_sha_none {
  address 192.168.69.179;
  tag-list tl1;
  address-mask 255.255.252.0;
  target-parameters TP2_v3_sha_none;
}
target-parameters TP2_v3_sha_none {
  parameters {
    message-processing-model v3;
    security-model usm;
    security-level none;
    security-name RU2_v3_sha_none;
  }
  notify-filter nf1;
}
notify N1_all_tl1_informs {
  type inform; # Replace inform with trap to convert informs to traps.
  tag tl1;
}
notify-filter nf1 {
  oid .1 include;
}
}
view all {
  oid .1 include;
}
}

```

You can convert the SNMPv3 informs to traps by setting the value of the **type** statement at the **[edit snmp v3 notify N1_all_tl1_informs]** hierarchy level to **trap** as shown in the following example:

Converting Informs to Traps

```
user@host# set snmp v3 notify N1_all_tl1_informs type trap
```

Configuring System Name, Location, Description, and Contact Information

Junos OS enables you to include the name and location of the system, administrative contact information, and a brief description of the system in the SNMP configuration.



NOTE: Always keep the name, location, contact, and description information configured and updated for all your devices that are managed by SNMP.

The following example shows a typical configuration.



TIP: Use quotation marks to enclose the system name, contact, location, and description information that contain spaces.

```
[edit]
snmp {
  name "snmp 001"; # Overrides the system name.
  contact "Juniper Berry, (650) 555 1234"; # Specifies the name and phone number of
    the administrator.
  location "row 11, rack C"; # Specifies the location of the device.
  description "M40 router with 8 FPCs" # Configures a description for the device.
}
```

Related Documentation

- [Understanding SNMP Implementation in Junos OS on page 15](#)
- [Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS on page 195](#)
- [Optimizing the Network Management System Configuration for the Best Results on page 81](#)
- [Configuring Options on Managed Devices for Better SNMP Response Time on page 83](#)
- [Managing Traps and Informs](#)
- [Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage](#)

Configuring the System Contact on a Device Running Junos OS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

You can specify an administrative contact for each system being managed by SNMP. This name is placed into the MIB II **sysContact** object. To configure a contact name, include the **contact** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
contact contact;
```

If the name contains spaces, enclose it in quotation marks (" ").

To define a system contact name that contains spaces:

```
[edit]
snmp {
  contact "Juniper Berry, (650) 555-1234";
}
```

- Related Documentation**
- [Configuring SNMP on a Device Running Junos OS](#)
 - [Configuring the System Location for a Device Running Junos OS on page 89](#)
 - [Configuring the System Description on a Device Running Junos OS on page 89](#)
 - [Configuring a Different System Name on page 91](#)
 - [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring the System Location for a Device Running Junos OS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

You can specify the location of each system being managed by SNMP. This string is placed into the MIB II **sysLocation** object. To configure a system location, include the **location** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
location location;
```

If the location contains spaces, enclose it in quotation marks (" ").

To specify the system location:

```
[edit]
snmp {
    location "Row 11, Rack C";
}
```

- Related Documentation**
- [Configuring SNMP on a Device Running Junos OS](#)
 - [Configuring the System Contact on a Device Running Junos OS on page 88](#)
 - [Configuring the System Description on a Device Running Junos OS on page 89](#)
 - [Configuring a Different System Name on page 91](#)
 - [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring the System Description on a Device Running Junos OS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

You can specify a description for each system being managed by SNMP. This string is placed into the MIB II **sysDescription** object. To configure a description, include the **description** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
description description;
```

If the description contains spaces, enclose it in quotation marks (" ").

To specify the system description:

```
[edit]
snmp {
```

```
description "M40 router with 8 FPCs";  
}
```

**Related
Documentation**

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuring the System Contact on a Device Running Junos OS on page 88](#)
- [Configuring the System Location for a Device Running Junos OS on page 89](#)
- [Configuring a Different System Name on page 91](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring SNMP Details

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

You can use SNMP to store basic administrative details, such as a contact name and the location of the device. Your management system can then retrieve this information remotely, when you are troubleshooting an issue or performing an audit. In SNMP terminology, these are the `sysContact`, `sysDescription`, and `sysLocation` objects found within the system group of MIB-2 (as defined in RFC 1213, *Management Information Base for Network Management of TCP/IP-based internets: MIB-II*). You can set initial values directly in the Junos OS configuration for each system being managed by SNMP.

To set the system contact details:

1. Set the system contact details by including the **contact** statement at the **[edit snmp]** hierarchy level, or in an appropriate configuration group as shown here.

This administrative contact is placed into the MIB II **sysContact** object.

If the name contains spaces, enclose it in quotation marks (" ").

```
[edit groups global snmp]  
user@host# set contact contact
```

For example:

```
[edit groups global snmp]  
user@host# set contact "Enterprise Support, (650) 555-1234"
```

2. Configure a system description.

This string is placed into the MIB II **sysDescription** object. If the description contains spaces, enclose it in quotation marks (" ").

```
[edit groups global snmp]  
user@host# set description description
```

For example:

```
[edit groups global snmp]  
user@host# set description "M10i router with 8 FPCs"
```

3. Configure a system location.

This string is placed into the MIB II **sysLocation** object. If the location contains spaces, enclose it in quotation marks (" ").

To specify the system location:

```
[edit]
snmp {
  location "Row 11, Rack C";
}

[edit groups global snmp]
user@host# set location location
```

For example:

```
[edit groups global snmp]
user@host# set location "London Corporate Office, Lab 5, Row 11, Rack C"
```

4. At the top level of the configuration, apply the configuration group.

If you use a configuration group, you must apply it for it to take effect.

```
[edit]
user@host# set apply-groups global
```

5. Commit the configuration.

```
user@host# commit
```

6. To verify the configuration, enter the **show snmp mib walk system** operational-mode command.

The **show snmp mib walk system** command performs a MIB walk through of the system table (from MIB-2 as defined in RFC 1213). The SNMP agent in Junos OS responds by printing each row in the table and its associated value. You can use the same command to perform a MIB walk through any part of the MIB tree supported by the agent.

```
user@host> show snmp mib walk system
sysDescr.0    = M10i router with 8 FPCs
sysObjectID.0 = jnxProductNameM10i
sysUpTime.0   = 173676474
sysContact.0  = Enterprise Support, (650) 555-1234
sysName.0     = host
sysLocation.0 = London Corporate Office, Lab 5, Row 11, Rack C
sysServices.0 = 4
```

Related Documentation

- [Configuring SNMP Communities on page 93](#)
- [Configuring SNMP Traps on page 100](#)
- [Configuring SNMP on a Device Running Junos OS](#)

Configuring a Different System Name

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

Junos OS enables you to override the system name by including the **name** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
name name;
```

If the name contains spaces, enclose it in quotation marks (" ").

To specify the system name override:

```
[edit]
snmp {
  name "snmp 1";
}
```

**Related
Documentation**

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuring the System Contact on a Device Running Junos OS on page 88](#)
- [Configuring the System Location for a Device Running Junos OS on page 89](#)
- [Configuring the System Description on a Device Running Junos OS on page 89](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring the Commit Delay Timer

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

When a router or switch first receives an SNMP nonvolatile **Set** request, a Junos OS XML protocol session opens and prevents other users or applications from changing the candidate configuration (equivalent to the command-line interface [CLI] **configure exclusive** command). If the router does not receive new SNMP **Set** requests within 5 seconds (the default value), the candidate configuration is committed and the Junos OS XML protocol session closes (the configuration lock is released). If the router receives new SNMP **Set** requests while the candidate configuration is being committed, the SNMP **Set** request is rejected and an error is generated. If the router receives new SNMP **Set** requests before 5 seconds have elapsed, the commit-delay timer (the length of time between when the last SNMP request is received and the commit is requested) resets to 5 seconds.

By default, the timer is set to 5 seconds. To configure the timer for the SNMP **Set** reply and start of the commit, include the **commit-delay** statement at the **[edit snmp nonvolatile]** hierarchy level:

```
[edit snmp nonvolatile]
commit-delay seconds;
```

seconds is the length of the time between when the SNMP request is received and the commit is requested for the candidate configuration. For more information about the **configure exclusive** command and locking the configuration, see the *CLI User Guide*.

**Related
Documentation**

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Filtering Duplicate SNMP Requests

Supported Platforms [PTX Series](#)

By default, filtering duplicate **get**, **getNext**, and **getBulk** SNMP requests is disabled on devices running Junos OS. If a network management station retransmits a **Get**, **GetNext**, or **GetBulk** SNMP request too frequently to the router, that request might interfere with the processing of previous requests and slow down the response time of the agent. Filtering these duplicate requests improves the response time of the SNMP agent. Junos OS uses the following information to determine if an SNMP request is a duplicate:

- Source IP address of the SNMP request
- Source UDP port of the SNMP request
- Request ID of the SNMP request

To filter duplicate SNMP requests, include the **filter-duplicates** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]  
filter-duplicates;
```

Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 109](#)
- [Filtering Interface Information Out of SNMP Get and GetNext Output on page 110](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring SNMP Communities

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Configuring the SNMP agent in Junos OS is a straightforward task that shares many familiar settings common to other managed devices in your network. For example, you need to configure Junos OS with an SNMP community string and a destination for traps. Community strings are administrative names that group collections of devices and the agents that are running on them together into common management domains. If a manager and an agent share the same community, they can communicate with each other. An SNMP community defines the level of authorization granted to its members, such as which MIB objects are available, which operations (read-only or read-write) are valid for those objects, and which SNMP clients are authorized, based on their source IP addresses.

The SNMP community string defines the relationship between an SNMP server system and the client systems. This string acts like a password to control the clients' access to the server.

To create a read-only SNMP community:

1. Enter the SNMP community used in your network.

If the community name contains spaces, enclose it in quotation marks (" ").

Community names must be unique.



NOTE: You cannot configure the same community name at the [edit snmp community] and [edit snmp v3 snmp-community *community-index*] hierarchy levels.

```
[edit groups global]
user@host# set snmp community name
```

This example uses the standard name **public** to create a community that gives limited read-only access.

```
[edit groups global]
user@host# set snmp community public
```

2. Define the authorization level for the community.

The default authorization level for a community is **read-only**.

To allow **Set** requests within a community, you need to define that community as **authorization read-write**. For **Set** requests, you also need to include the specific MIB objects that are accessible with read-write privileges using the **view** statement. The default view includes all supported MIB objects that are accessible with read-only privileges. No MIB objects are accessible with read-write privileges. For more information about the **view** statement, see [“Configuring MIB Views” on page 111](#).

```
[edit groups global snmp community name]
user@host# set authorization authorization
```

This example confines the public community to read-only access. Any SNMP client (for example, an SNMP management system) that belongs to the public community can read MIB variables but cannot set (change) them.

```
[edit groups global snmp community public]
user@host# set authorization read-only
```

3. Define a list of clients in the community who are authorized to communicate with the SNMP agent in Junos OS.

The **clients** statement lists the IP addresses of the clients (community members) that are allowed to use this community. List the clients by IP address and prefix. Typically, the list includes the SNMP network management system in your network or the address of your management network. If no **clients** statement is present, all clients are allowed. For **address**, you must specify an IPv4 or IPv6 address, not a hostname.

```
[edit groups global snmp community name]
user@host# set clients address
```


The following statement defines the hosts in the 192.168.1.0/24 network as being authorized in the public community.

```
[edit groups global snmp community public]
user@host# set clients 192.168.1.0/24
```

4. Define the clients that are not authorized within the community by specifying their IP address, followed by the **restrict** statement.

```
[edit groups global snmp community name]
user@host# set clients address restrict
```

The following statement defines all other hosts as being restricted from the public community.

```
[edit groups global snmp community public]
user@host# set clients 0/0 restrict
```

5. At the top level of the configuration, apply the configuration group.

If you use a configuration group, you must apply it for it to take effect.

```
[edit]
user@host# set apply-groups global
```

6. Commit the configuration.

```
user@host# commit
```

To create a read-write SNMP community:

1. Enter the SNMP community used in your network.

```
[edit groups global]
user@host# set snmp community name
```

This example standard community string **private** to identify the community granted read-write access to the SNMP agent running on the device.

```
[edit groups global]
user@host# set snmp community private
```

2. Define the authorization level for the community.

```
[edit groups global snmp community name]
user@host# set authorization authorization
```

This example confines the public community to read-only access. Any SNMP client (for example, an SNMP management system) that belongs to the public community can read MIB variables but cannot set (change) them.

```
[edit groups global snmp community public]
user@host# set authorization read-write
```

3. Define a list of clients in the community who are authorized to make changes to the SNMP agent in Junos OS.

List the clients by IP address and prefix.

```
[edit groups global snmp community name]  
user@host# set clients address
```

For example:

```
[edit groups global snmp community private]  
user@host# set clients 192.168.1.15/24  
user@host# set clients 192.168.1.18/24
```

4. Define the clients that are not authorized within the community by specifying their IP address, followed by the **restrict** statement.

```
[edit groups global snmp community name]  
user@host# set clients address restrict
```

The following statement defines all other hosts as being restricted from the public community.

```
[edit groups global snmp community private]  
user@host# set clients 0/0 restrict
```

5. At the top level of the configuration, apply the configuration group.

If you use a configuration group, you must apply it for it to take effect.

```
[edit]  
user@host# set apply-groups global
```

6. Commit the configuration.

```
user@host# commit
```

Related Documentation

- [Adding a Group of Clients to an SNMP Community on page 97](#)
- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Examples: Configuring the SNMP Community String on page 96](#)

Examples: Configuring the SNMP Community String

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Grant read-only access to all clients. With the following configuration, the system responds to SNMP **Get**, **GetNext**, and **GetBulk** requests that contain the community string **public**:

```
[edit]  
snmp {  
  community public {  
    authorization read-only;  
  }  
}
```

Grant all clients read-write access to the ping MIB and **jnxPingMIB**. With the following configuration, the system responds to SNMP **Get**, **GetNext**, **GetBulk**, and **Set** requests that contain the community string **private** and specify an OID contained in the ping MIB or **jnxPingMIB** hierarchy:

```
[edit]
snmp {
  view ping-mib-view {
    oid pingMIB include;
    oid jnxPingMIB include;
    community private {
      authorization read-write;
      view ping-mib-view;
    }
  }
}
```

The following configuration allows read-only access to clients with IP addresses in the range 1.2.3.4/24, and denies access to systems in the range fe80::1:2:3:4/64:

```
[edit]
snmp {
  community field-service {
    authorization read-only;
    clients {
      default restrict; # Restrict access to all SNMP clients not explicitly
                        # listed on the following lines.
      1.2.3.4/24; # Allow access by all clients in 1.2.3.4/24 except
      fe80::1:2:3:4/64 restrict;# fe80::1:2:3:4/64.
    }
  }
}
```

Related Documentation • [Configuring SNMP Communities on page 93](#)

Adding a Group of Clients to an SNMP Community

Supported Platforms ACX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX

Junos OS enables you to add one or more groups of clients to an SNMP community. You can include the **client-list-name** *name* statement at the **[edit snmp community community-name]** hierarchy level to add all the members of the client list or prefix list to an SNMP community.

To define a list of clients, include the **client-list** statement followed by the IP addresses of the clients at the **[edit snmp]** hierarchy level:

```
[edit snmp]
  client-list client-list-name {
    ip-addresses;
  }
```

You can configure a prefix list at the **[edit policy options]** hierarchy level. Support for prefix lists in the SNMP community configuration enables you to use a single list to

configure the SNMP and routing policies. For more information about the **prefix-list** statement, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

To add a client list or prefix list to an SNMP community, include the **client-list-name** statement at the **[edit snmp community *community-name*]** hierarchy level:

```
[edit snmp community community-name]  
client-list-name client-list-name;
```



NOTE: The client list and prefix list must not have the same name.

The following example shows how to define a client list:

```
[edit]  
snmp {  
  client-list clentlist1 {  
    10.1.1.1/32;  
    10.2.2.2/32;  
  }  
}
```

The following example shows how to add a client list to an SNMP community:

```
[edit]  
snmp {  
  community community1 {  
    authorization read-only;  
    client-list-name clientlist1;  
  }  
}
```

The following example shows how to add a prefix list to an SNMP community:

```
[edit]  
policy-options {  
  prefix-list prefixlist {  
    10.3.3.3/32;  
    10.5.5.5/32;  
  }  
}  
snmp {  
  community community2 {  
    client-list-name prefixlist;  
  }  
}
```

Related Documentation

- *client-list*
- *client-list-name*

Configuring a Proxy SNMP Agent

Supported Platforms [M Series](#), [MX Series](#), [T Series](#)

Starting with Release 12.3, Junos OS enables you to assign one of the devices in the network as a proxy SNMP agent through which the network management system (NMS) can query other devices in the network. When you configure a proxy, you can specify the names of devices to be managed through the proxy SNMP agent.

When the NMS queries the proxy SNMP agent, the NMS specifies the community name (for SNMPv1 and SNMPv2) or the context and security name (for SNMPv3) associated with the device from which it requires the information.



NOTE: If you have configured authentication and privacy methods and passwords for SNMPv3, those parameters are also specified in the query for SNMPv3 information.

To configure a proxy SNMP agent and specify devices to be managed by the proxy SNMP agent, you can include the following configuration statements at the `[edit snmp]` hierarchy level:

```
proxy proxy-name{
  device-name device-name;
  logical-system logical-system {
    routing-instance routing-instance;
  }
  routing-instance routing-instance;
  <version-v1 | version-v2c> {
    snmp-community community-name;
    no-default-comm-to-v3-config;
  }
  version-v3 {
    security-name security-name;
    context context-name;
  }
}
```

- The **proxy** statement enables you to specify a unique name for the proxy configuration.
- The **version-v1**, **version-v2c**, and **version-v3** statements enable you to specify the SNMP version.
- The **no-default-comm-to-v3-config** statement is an optional statement at the `[edit snmp proxy proxy-name <version-v1 | version-v2c>]` hierarchy level that when included in the configuration requires you to manually configure the statements at the `[edit snmp v3 snmp-community community-name]` and `[edit snmp v3 vacm]` hierarchy levels.

If the **no-default-comm-to-v3-config** statement is not included at the `[edit snmp proxy proxy-name <version-v1 | version-v2c>]` hierarchy level, the `[edit snmp v3 snmp-community community-name]` and `[edit snmp v3 vacm]` hierarchy level configurations are automatically initialized.

- The **logical-system** and **routing-instance** statements are optional statements that enable you to specify logical system and routing instance names if you want to create proxies for logical systems or routing instances on the device.



NOTE: Starting with Junos OS Release 15.2, you must configure interface `<interface-name>` statement at the `[edit snmp]` hierarchy level for the proxy SNMP agent.



NOTE: The community and security configuration for the proxy should match the corresponding configuration on the device that is to be managed.



NOTE: Because the proxy SNMP agent does not have trap forwarding capabilities, the devices that are managed by the proxy SNMP agent send the traps directly to the network management system.

You can use the **show snmp proxy** operational mode command to view proxy details on a device. The **show snmp proxy** command returns the proxy names, device names, SNMP version, community/security, and context information.

Related Documentation

- [proxy \(snmp\) on page 750](#)

Configuring SNMP Traps

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Traps are unsolicited messages sent from an SNMP agent to remote network management systems or trap receivers. Many enterprises use SNMP traps as part of a fault-monitoring solution, in addition to system logging. In Junos OS, SNMP traps are not forwarded by default, so you must configure a trap-group if you wish to use SNMP traps.

You can create and name a group of one or more types of SNMP traps and then define which systems receive the group of SNMP traps.. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name.

To configure an SNMP trap:

1. Create a single, consistent source address that Junos OS applies to all outgoing traps in your device.

A source address is useful, because although most Junos OS devices have a number of outbound interfaces, using one source address helps a remote NMS to associate the source of the traps with an individual device

```
[edit groups global snmp]
user@host# set trap-options source-address address
```

This example uses the IP address of the loopback interface (lo0) as the source address for all the SNMP traps that originate from the device.

```
[edit groups global snmp]
```

```
user@host# set trap-options source-address lo0
```

2. Create a trap group in which you can list the types of traps to be forwarded and the targets (addresses) of the receiving remote management systems.

```
[edit groups global snmp trap-group group-name]
user@host# set version (all | v1 | v2) targets address
```

This example creates a trap group called **managers**, allows SNMP version 2-formatted notifications (traps) to be sent to the host at address 192.168.1.15. This statement forwards all categories of traps.

```
[edit groups global snmp trap-group managers]
user@host# set version v2 targets 192.168.1.15
```

3. Define the specific subset of trap categories to be forwarded.

For a list of categories, see [“Configuring SNMP Trap Groups” on page 106](#).

```
[edit groups global snmp trap-group group-name]
user@host# set categories category
```

The following statement configures the standard MIB-II authentication failures on the agent (the device).

```
[edit groups global snmp trap-group managers]
user@host# set categories authentication
```

4. At the top level of the configuration, apply the configuration group.

If you use a configuration group, you must apply it for it to take effect.

```
[edit]
user@host# set apply-groups global
```

5. Commit the configuration.

```
user@host# commit
```

6. To verify the configuration, generate an authentication failure trap.

This means that the SNMP agent received a request with an unknown community. Other traps types can also be spoofed as well.

This feature enables you to trigger SNMP traps from routers and ensure that they are processed correctly within your existing network management infrastructure. This is also useful for testing and debugging SNMP behavior on the switch or NMS.

Using the **monitor traffic** command, you can verify that the trap is sent to the network management system.

```
user@host> request snmp spoof-trap authenticationFailure
Spoof-trap request result: trap sent successfully
```

Related Documentation

- [Adding a Group of Clients to an SNMP Community on page 97](#)
- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Examples: Configuring the SNMP Community String on page 96](#)

Configuring SNMP Trap Options and Groups on a Device Running Junos OS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Some carriers have more than one trap receiver that forwards traps to a central NMS. This allows for more than one path for SNMP traps from a router to the central NMS through different trap receivers. A device running Junos OS can be configured to send the same copy of each SNMP trap to every trap receiver configured in the trap group.

The source address in the IP header of each SNMP trap packet is set to the address of the outgoing interface by default. When a trap receiver forwards the packet to the central NMS, the source address is preserved. The central NMS, looking only at the source address of each SNMP trap packet, assumes that each SNMP trap came from a different source.

In reality, the SNMP traps came from the same router, but each left the router through a different outgoing interface.

The statements discussed in the following sections are provided to allow the NMS to recognize the duplicate traps and to distinguish SNMPv1 traps based on the outgoing interface.

To configure SNMP trap options and trap groups, include the **trap-options** and **trap-group** statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  targets {
    address;
  }
  version (all | v1 | v2);
}
```

Related Documentation

- [Configuring SNMP Trap Options on page 103](#)
- [Configuring SNMP Trap Groups on page 106](#)
- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring SNMP Trap Options

Supported Platforms M Series, MX Series, PTX Series, T Series

Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router to a single address regardless of the outgoing interface. In addition, you can set the agent address of the SNMPv1 traps. For more information about the contents of SNMPv1 traps, see RFC 1157.



NOTE: SNMP cannot be associated with any routing instances other than the master routing instance.

To configure SNMP trap options, include the **trap-options** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
trap-options {
  agent-address outgoing-interface;
  enterprise-oid
  logical-system
  routing-instance
  source-address address;
}
```

You must also configure a trap group for the trap options to take effect. For information about trap groups, see “Configuring SNMP Trap Groups” on page 106.

This topic contains the following sections:

- [Configuring the Source Address for SNMP Traps on page 103](#)
- [Configuring the Agent Address for SNMP Traps on page 105](#)
- [Adding snmpTrapEnterprise Object Identifier to Standard SNMP Traps on page 106](#)

Configuring the Source Address for SNMP Traps

You can configure the source address of trap packets in many ways: **lo0**, a valid IPv4 address or IPv6 address configured on one of the router interfaces, a logical-system address, or the address of a routing-instance. The value **lo0** indicates that the source address of the SNMP trap packets is set to the lowest loopback address configured on the interface **lo0**.



NOTE: If the source address is an invalid IPv4 or IPv6 address or is not configured, SNMP traps are not generated.

You can configure the source address of trap packets in one of the following formats:

- A valid IPv4 address configured on one of the router interfaces
- A valid IPv6 address configured on one of the router interfaces

- **lo0**; that is, the lowest loopback address configured on the interface **lo0**
- A logical-system name
- A routing-instance name

A Valid IPv4 Address As the Source Address

To specify a valid IPv4 interface address as the source address for SNMP traps on one of the router interfaces, include the **source-address** statement at the **[edit snmp trap-options]** hierarchy level:

```
[edit snmp trap-options]
source-address address;
```

address is a valid IPv4 address configured on one of the router interfaces.

A Valid IPv6 Address As the Source Address

To specify a valid IPv6 interface address as the source address for SNMP traps on one of the router interfaces, include the **source-address** statement at the **[edit snmp trap-options]** hierarchy level:

```
[edit snmp trap-options]
source-address address;
```

address is a valid IPv6 address configured on one of the router interfaces.

The Lowest Loopback Address As the Source Address

To specify the source address of the SNMP traps so that they use the lowest loopback address configured on the interface **lo0** as the source address, include the **source-address** statement at the **[edit snmp trap-options]** hierarchy level:

```
[edit snmp trap-options]
source-address lo0;
```

To enable and configure the loopback address, include the **address** statement at the **[edit interfaces lo0 unit 0 family inet]** hierarchy level:

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address ip-address;
    }
  }
}
```

To configure the loopback address as the source address of trap packets:

```
[edit snmp]
trap-options {
  source-address lo0;
}
trap-group "urgent-dispatcher" {
  version v2;
  categories link startup;
  targets {
    192.168.10.22;
    172.17.1.2;
  }
}
```

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address 10.0.0.1/32;
      address 127.0.0.1/32;
    }
  }
}
```

In this example, the IP address **10.0.0.1** is the source address of every trap sent from this router.

Logical System Name as the Source Address

To specify a logical system name as the source address of SNMP traps, include the **logical-system** *logical-system-name* statement at the **[edit snmp trap-options]** hierarchy level.

For example, the following configuration sets logical system name **ls1** as the source address of SNMP traps:

```
[edit snmp]
  trap-options {
    logical-system ls1;
  }
```

Routing Instance Name as the Source Address

To specify a routing instance name as the source address of SNMP traps, include the **routing-instance** *routing-instance-name* statement at the **[edit snmp trap-options]** hierarchy level.

For example, the following configuration sets the routing instance name **ri1** as the source address for SNMP traps:

```
[edit snmp]
  trap-options {
    routing-instance ri1;
  }
```

Configuring the Agent Address for SNMP Traps

The agent address is only available in SNMPv1 trap packets (see RFC 1157). By default, the router's default local address is not specified in the agent address field of the SNMPv1 trap. To configure the agent address, include the **agent-address** statement at the **[edit snmp trap-options]** hierarchy level. Currently, the agent address can only be the address of the outgoing interface:

```
[edit snmp]
  trap-options {
    agent-address outgoing-interface;
  }
```

To configure the outgoing interface as the agent address:

```
[edit snmp]
  trap-options {
    agent-address outgoing-interface;
  }
```

```
trap-group "urgent-dispatcher" {  
  version v1;  
  categories link startup;  
  targets {  
    192.168.10.22;  
    172.17.1.2;  
  }  
}
```

In this example, each SNMPv1 trap packet sent has its agent address value set to the IP address of the outgoing interface.

Adding snmpTrapEnterprise Object Identifier to Standard SNMP Traps

The **snmpTrapEnterprise** object helps you identify the enterprise that has defined the trap. Typically, the **snmpTrapEnterprise** object appears as the last varbind in enterprise-specific SNMP version 2 traps. However, starting Release 10.0, Junos OS enables you to add the **snmpTrapEnterprise** object identifier to standard SNMP traps as well.

To add **snmpTrapEnterprise** to standard traps, include the **enterprise-oid** statement at the **[edit snmp trap-options]** hierarchy level. If the **enterprise-oid** statement is not included in the configuration, **snmpTrapEnterprise** is added only for enterprise-specific traps.

```
[edit snmp]  
trap-options {  
  enterprise-oid;  
}
```

Related Documentation

- [Configuring SNMP Trap Options and Groups on a Device Running Junos OS on page 102](#)
- [Configuring SNMP Trap Groups on page 106](#)
- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring SNMP Trap Groups

Supported Platforms [SRX Series, vSRX](#)

You can create and name a group of one or more types of SNMP traps and then define which systems receive the group of SNMP traps. The trap group must be configured for SNMP traps to be sent. To create an SNMP trap group, include the **trap-group** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]  
trap-group group-name {  
  categories {  
    category;  
  }  
  destination-port port-number;  
  routing-instance instance;  
  targets {  
    address;
```

```

    }
    version (all | v1 | v2);
  }

```

The trap group name can be any string and is embedded in the community name field of the trap. To configure your own trap group port, include the **destination-port** statement. The default destination port is port 162.

For each trap group that you define, you must include the **target** statement to define at least one system as the recipient of the SNMP traps in the trap group. Specify the IPv4 or IPv6 address of each recipient, not its hostname.

Specify the types of traps the trap group can receive in the **categories** statement. For information about the category to which the traps belong, see the “[Standard SNMP Traps Supported by Junos OS](#)” on page 50 and “[Enterprise-Specific SNMP Traps Supported by Junos OS](#)” on page 57 topics.

Specify the routing instance used by the trap group in the **routing-instance** statement. All targets configured in the trap group use this routing instance.

A trap group can receive the following categories:

- **authentication**—Authentication failures
- **chassis**—Chassis or environment notifications
- **configuration**—Configuration notifications
- **link**—Link-related notifications (up-down transitions, DS-3 and DS-1 line status change, IPv6 interface state change, and Passive Monitoring PIC overload)



NOTE: To send Passive Monitoring PIC overload interface traps, select the link trap category.

- **remote-operations**—Remote operation notifications
- **rmon-alarm**—Alarm for RMON events
- **routing**—Routing protocol notifications
- **sonet-alarms**—SONET/SDH alarms



NOTE: If you omit the SONET/SDH subcategories, all SONET/SDH trap alarm types are included in trap notifications.

- **loss-of-light**—Loss of light alarm notification
- **pll-lock**—PLL lock alarm notification
- **loss-of-frame**—Loss of frame alarm notification
- **loss-of-signal**—Loss of signal alarm notification
- **severely-errored-frame**—Severely errored frame alarm notification

- **line-ais**—Line alarm indication signal (AIS) alarm notification
- **path-ais**—Path AIS alarm notification
- **loss-of-pointer**—Loss of pointer alarm notification
- **ber-defect**—SONET/SDH bit error rate alarm defect notification
- **ber-fault**—SONET/SDH error rate alarm fault notification
- **line-remote-defect-indication**—Line remote defect indication alarm notification
- **path-remote-defect-indication**—Path remote defect indication alarm notification
- **remote-error-indication**—Remote error indication alarm notification
- **unequipped**—Unequipped alarm notification
- **path-mismatch**—Path mismatch alarm notification
- **loss-of-cell**—Loss of cell delineation alarm notification
- **vt-ais**—Virtual tributary (VT) AIS alarm notification
- **vt-loss-of-pointer**—VT loss of pointer alarm notification
- **vt-remote-defect-indication**—VT remote defect indication alarm notification
- **vt-unequipped**—VT unequipped alarm notification
- **vt-label-mismatch**—VT label mismatch error notification
- **vt-loss-of-cell**—VT loss of cell delineation notification
- **startup**—System warm and cold starts
- **timing-events**—Timing events and defects notification
- **vrrp-events**—Virtual Router Redundancy Protocol (VRRP) events such as new-master or authentication failures
- **startup**—System warm and cold starts
- **vrrp-events**—Virtual Router Redundancy Protocol (VRRP) events such as new-master or authentication failures

If you include SONET/SDH subcategories, only those SONET/SDH trap alarm types are included in trap notifications.

The **version** statement allows you to specify the SNMP version of the traps sent to targets of the trap group. If you specify **v1** only, SNMPv1 traps are sent. If you specify **v2** only, SNMPv2 traps are sent. If you specify **all**, both an SNMPv1 and an SNMPv2 trap are sent for every trap condition. For more information about the **version** statement, see [version \(SNMP\)](#).

Related Documentation

- [Configuring SNMP Trap Options and Groups on a Device Running Junos OS on page 102](#)
- [Configuring SNMP Trap Options on page 103](#)
- [Configuring SNMP on a Device Running Junos OS](#)

- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Example: Configuring SNMP Trap Groups on page 109](#)

Example: Configuring SNMP Trap Groups

Supported Platforms [M Series, MX Series, PTX Series, T Series](#)

Set up a trap notification list named **urgent-dispatcher** for link and startup traps. This list is used to identify the network management hosts (**1.2.3.4** and **fe80::1:2:3:4**) to which traps generated by the local router should be sent. The name specified for a trap group is used as the SNMP community string when the agent sends traps to the listed targets.

```
[edit]
snmp {
  trap-group "urgent-dispatcher" {
    version v2;
    categories link startup;
    targets {
      1.2.3.4;
      fe80::1:2:3:4;
    }
  }
}
```

- Related Documentation**
- [Configuring SNMP Trap Groups on page 106](#)
 - [Configuring SNMP Trap Options and Groups on a Device Running Junos OS on page 102](#)
 - [Configuring SNMP Trap Options on page 103](#)

Configuring the Interfaces on Which SNMP Requests Can Be Accepted

Supported Platforms [M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series](#)

By default, all router or switch interfaces have SNMP access privileges. To limit the access through certain interfaces only, include the **interface** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
interface [ interface-names ];
```

Specify the names of any logical or physical interfaces that should have SNMP access privileges. Any SNMP requests entering the router or switch from interfaces not listed are discarded.

- Related Documentation**
- [Configuring SNMP on a Device Running Junos OS](#)
 - [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
 - [Example: Configuring Secured Access List Checking on page 110](#)

Example: Configuring Secured Access List Checking

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

SNMP access privileges are granted to only devices on interfaces **so-0/0/0** and **at-1/0/1**. The following example does this by configuring a list of logical interfaces:

```
[edit]
snmp {
  interface [ so-0/0/0.0 so-0/0/0.1 at-1/0/1.0 at-1/0/1.1 ];
}
```

The following example grants the same access by configuring a list of physical interfaces:

```
[edit]
snmp {
  interface [ so-0/0/0 at-1/0/1 ];
}
```

- Related Documentation**
- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 109](#)
 - [Filtering Interface Information Out of SNMP Get and GetNext Output on page 110](#)
 - [Configuring SNMP on a Device Running Junos OS](#)
 - [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Filtering Interface Information Out of SNMP Get and GetNext Output

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Junos OS enables you to filter out information related to specific interfaces from the output of SNMP **Get** and **GetNext** requests performed on interface-related MIBs such as IF MIB, ATM MIB, RMON MIB, and the Juniper Networks enterprise-specific IF MIB.

You can use the following options of the **filter-interfaces** statement at the **[edit snmp]** hierarchy level to specify the interfaces that you want to exclude from SNMP **Get** and **GetNext** queries:

- **interfaces**—Interfaces that match the specified regular expressions.
- **all-internal-interfaces**—Internal interfaces.

```
[edit]
snmp {
  filter-interfaces {
    interfaces {
      interface-name 1;
      interface-name 2;
    }
    all-internal-interfaces;
  }
}
```


Starting with Release 12.1, Junos OS provides an except option (! operator) that enables you to filter out all interfaces except those interfaces that match all the regular expressions prefixed with the ! mark.

For example, to filter out all interfaces except the **ge** interfaces from the SNMP **get** and **get-next** results, enter the following command:

```
[edit snmp]
user@host# set filter-interfaces interfaces "!~ge-.*"
user@host# commit
```

When this is configured, Junos OS filters out all interfaces except the **ge** interfaces from the SNMP **get** and **get-next** results.



NOTE: The ! mark is supported only as the first character of the regular expression. If it appears anywhere else in a regular expression, Junos OS considers the regular expression invalid, and returns an error.

However, note that these settings are limited to SNMP operations, and the users can continue to access information related to the interfaces (including those hidden using the **filter-interfaces** options) using the appropriate Junos OS command-line interface (CLI) commands.

Related Documentation

- [Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 109](#)
- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

Configuring MIB Views

Supported Platforms ACX Series, EX4600, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX

SNMPv3 defines the concept of MIB views in RFC 3415, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*. MIB views provide an agent better control over who can access specific branches and objects within its MIB tree. A view consists of a name and a collection of SNMP object identifiers, which are either explicitly included or excluded. Once defined, a view is then assigned to an SNMPv3 group or SNMPv1/v2c community (or multiple communities), automatically masking which parts of the agent's MIB tree members of the group or community can (or cannot) access.

By default, an SNMP community grants read access and denies write access to all supported MIB objects (even communities configured as **authorization read-write**). To restrict or grant read or write access to a set of MIB objects, you must configure a MIB view and associate the view with a community.

To configure MIB views, include the **view** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
view view-name {
```

```
oid object-identifier (include | exclude);  
}
```

The **view** statement defines a MIB view and identifies a group of MIB objects. Each MIB object of a view has a common object identifier (OID) prefix. Each object identifier represents a subtree of the MIB object hierarchy. The subtree can be represented either by a sequence of dotted integers (such as **1.3.6.1.2.1.2**) or by its subtree name (such as **interfaces**). A configuration statement uses a view to specify a group of MIB objects on which to define access. You can also use a wildcard character asterisk (*) to include OIDs that match a particular pattern in the SNMP view. To enable a view, you must associate the view with a community.

To remove an OID completely, use the **delete view all oid oid-number** command but omit the **include** parameter.

```
[edit groups global snmp]  
user@host# set view view-name oid object-identifier (include | exclude)
```

The following example creates a MIB view called ping-mib-view. The **oid** statement does not require a dot at the beginning of the object identifier. The **snmp view** statement includes the branch under the object identifier .1.3.6.1.2.1.80. This includes the entire DISMAN-PINGMIB subtree (as defined in RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*), which effectively permits access to any object under that branch.

```
[edit groups global snmp]  
user@host# set view ping-mib-view oid 1.3.6.1.2.1.80 include
```

The following example adds a second branch in the same MIB view.

```
[edit groups global snmp]  
user@host# set view ping-mib-view oid jnxPingMIB include
```

Assign a MIB view to a community that you want to control.

To associate MIB views with a community, include the **view** statement at the **[edit snmp community community-name]** hierarchy level:

```
[edit snmp community community-name]  
view view-name;
```

For more information about the Ping MIB, see RFC 2925 and [PING MIB](#).

Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Configuring Ping Proxy MIB on page 113](#)
- [view \(Configuring a MIB View\) on page 763](#)
- [view \(Associating MIB View with a Community\)](#)
- [oid on page 749](#)

Configuring Ping Proxy MIB

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Restrict the **ping-mib** community to read and write access of the Ping MIB and **jnxpingMIB** only. Read or write access to any other MIB using this community is not allowed.

```
[edit snmp]
view ping-mib-view {
  oid 1.3.6.1.2.1.80 include; #pingMIB
  oid jnxPingMIB include; #jnxPingMIB
}
community ping-mib {
  authorization read-write;
  view ping-mib-view;
}
```

The following configuration prevents the **no-ping-mib** community from accessing Ping MIB and **jnxPingMIB** objects. However, this configuration does not prevent the **no-ping-mib** community from accessing any other MIB object that is supported on the device.

```
[edit snmp]
view no-ping-mib-view {
  oid 1.3.6.1.2.1.80 exclude; # deny access to pingMIB objects
  oid jnxPingMIB exclude; # deny access to jnxPingMIB objects
}
community no-ping-mib {
  authorization read-write;
  view ping-mib-view;
}
```

**Related
Documentation**

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Configuring MIB Views on page 111](#)
- [view \(Configuring a MIB View\) on page 763](#)
- [oid on page 749](#)

CHAPTER 7

Configuring SNMPv3

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
- [Example: SNMPv3 Configuration on page 117](#)
- [Configuring the Local Engine ID on page 120](#)
- [Creating SNMPv3 Users on page 121](#)
- [Example: Creating SNMPv3 Users on page 122](#)
- [Configuring the SNMPv3 Authentication Type on page 123](#)
- [Configuring the SNMPv3 Encryption Type on page 124](#)
- [Defining Access Privileges for an SNMP Group on page 126](#)
- [Configuring the Access Privileges Granted to a Group on page 127](#)
- [Example: Configuring the Access Privileges Granted to a Group on page 130](#)
- [Assigning Security Model and Security Name to a Group on page 131](#)
- [Example: Security Group Configuration on page 133](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring the SNMPv3 Trap Notification on page 135](#)
- [Example: Configuring SNMPv3 Trap Notification on page 136](#)
- [Configuring the Trap Notification Filter on page 136](#)
- [Configuring the Trap Target Address on page 137](#)
- [Example: Configuring the Tag List on page 140](#)
- [Defining and Configuring the Trap Target Parameters on page 140](#)
- [Configuring SNMP Informs on page 143](#)
- [Configuring the Remote Engine and Remote User on page 144](#)
- [Example: Configuring the Remote Engine ID and Remote User on page 145](#)
- [Configuring the Inform Notification Type and Target Address on page 148](#)
- [Example: Configuring the Inform Notification Type and Target Address on page 150](#)
- [Configuring the SNMPv3 Community on page 151](#)
- [Example: Configuring an SNMPv3 Community on page 153](#)

Minimum SNMPv3 Configuration on a Device Running Junos OS

Supported Platforms ACX Series, EX4600, M Series, MX Series, PTX Series, QFabric System, QFX Series, T Series

To configure the minimum requirements for SNMPv3, include the following statements at the `[edit snmp v3]` and `[edit snmp]` hierarchy levels:



NOTE: You must configure at least one view (notify, read, or write) at the `[edit snmp view-name]` hierarchy level.

```
[edit snmp]
view view-name {
  oid object-identifier (include | exclude);
}
[edit snmp v3]
notify name {
  tag tag-name;
}
notify-filter profile-name {
  oid object-identifier (include | exclude);
}
snmp-community community-index {
  security-name security-name;
}
target-address target-address-name {
  address address;
  target-parameters target-parameters-name;
}
target-parameters target-parameters-name {
  notify-filter profile-name;
  parameters {
    message-processing-model (v1 | v2c | v3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
    }
  }
}
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix){
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
}
}

```

Related Documentation

- [Creating SNMPv3 Users on page 121](#)
- [Configuring MIB Views on page 111](#)
- [Defining Access Privileges for an SNMP Group on page 126](#)
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring SNMP Informs on page 143](#)
- [Example: SNMPv3 Configuration on page 117](#)

Example: SNMPv3 Configuration

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Define an SNMPv3 configuration:

```

[edit snmp]
engine-id {
  use-mac-address;
}
view jnxAlarms {
  oid 1.3.6.1.4.1.2636.3.4 include;
}
view interfaces {
  oid 1.3.6.1.2.1.2 include;
}
view ping-mib {
  oid 1.3.6.1.2.1.80 include;
}
[edit snmp v3]
notify n1 {
  tag router1; # Identifies a set of target addresses
  type trap; # Defines type of notification
}
notify n2 {
  tag host1;
  type trap;
}
notify-filter nf1 {
  oid .1 include; # Defines which traps to send
} # In this case, includes all traps

```

```
notify-filter nf2 {
  oid 1.3.6.1.4.1 include; # Sends enterprise-specific traps only
}
notify-filter nf3 {
  oid 1.3.6.1.2.1.1.5 include; # Sends BGP traps only
}
snmp-community index1 {
  community-name "$9$JOZi.QF/AtOz3"; # SECRET-DATA
  security-name john; # Matches the security name at the target parameters
  tag host1; # Finds the addresses that are allowed to be used with
}
target-address ta1 { # Associates the target address with the group
  # san-francisco.
  address 10.1.1.1;
  address-mask 255.255.255.0; # Defines the range of addresses
  port 162;
  tag-list router1;
  target-parameters tp1; # Applies configured target parameters
}
target-address ta2 {
  address 10.1.1.2;
  address-mask 255.255.255.0;
  port 162;
  tag-list host1;
  target-parameters tp2;
}
target-address ta3 {
  address 10.1.1.3;
  address-mask 255.255.255.0;
  port 162;
  tag-list "router1 host1";
  target-parameters tp3;
}
target-parameters tp1 { # Defines the target parameters
  notify-filter nf1; # Specifies which notify filter to apply
  parameters {
    message-processing-model v1;
    security-model v1;
    security-level none;
    security-name john; # Matches the security name configured at the
  } # [edit snmp v3 snmp-community community-index hierarchy level.
}
target-parameters tp2 {
  notify-filter nf2;
  parameters {
    message-processing-model v1;
    security-model v1;
    security-level none;
    security-name john;
  }
}
target-parameters tp3 {
  notify-filter nf3;
  parameters {
    message-processing-model v1;
    security-model v1;
```



```

        security-level none;
        security-name john;
    }
}
usm {
    local-engine { # Defines authentication and encryption for SNMPv3 users
        user john { # security-name john is defined here
            authentication-md5 {
                authentication-password authentication-password;
            }
            privacy-des {
                privacy-password privacy-password;
            }
        }
        user bob { # security-name bob is defined here
            authentication-sha {
                authentication-password authentication-password;
            }
            privacy-none;
        }
        user julia { # security-name julia is defined here
            authentication-none;
            privacy-none;
        }
        user lauren { # security-name lauren is defined here
            authentication-sha {
                authentication-password authentication-password;
            }
            privacy-aes128 {
                privacy-password privacy-password;
            }
        }
        user richard { # security-name richard is defined here
            authentication-sha {
                authentication-password authentication-password;
            }
            privacy-none;
        }
    }
}
vacm {
    access {
        group san-francisco { #Defines the access privileges for the group
            default-context-prefix { # called san-francisco
                security-model v1 {
                    security-level none {
                        notify-view ping-mib;
                        read-view interfaces;
                        write-view jnxAlarms;
                    }
                }
            }
        }
    }
}
security-to-group {
    security-model v1 {

```

```
security-name john { # Assigns john to security group san-francisco
  group san-francisco;
}
security-name bob { # Assigns bob to security group new-york
  group new-york;
}
security-name julia { # Assigns julia to security group chicago
  group chicago;
}
security-name lauren { # Assigns lauren to security group paris
  group paris;
}
security-name richard { # Assigns richard to security group geneva
  group geneva;
}
}
}
```

Related Documentation

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the Local Engine ID

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

By default, the local engine ID uses the default IP address of the router. The local engine ID is the administratively unique identifier for the SNMPv3 engine. This statement is optional. To configure the local engine ID, include the **engine-id** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
engine-id {
  (local engine-id-suffix | use-default-ip-address | use-mac-address);
}
```

- **local engine-id-suffix**—The engine ID suffix is explicitly configured.
- **use-default-ip-address**—The engine ID suffix is generated from the default IP address.
- **use-mac-address**—The SNMP engine identifier is generated from the Media Access Control (MAC) address of the management interface on the router.

The local engine ID is defined as the administratively unique identifier of an SNMPv3 engine, and is used for identification, not for addressing. There are two parts of an engine ID: prefix and suffix. The prefix is formatted according to the specifications defined in RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*. You can configure the suffix here.



NOTE: SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID. If you configure or change the engine ID, you must commit the new engine ID before you configure SNMPv3 users. Otherwise the keys generated from the configured passwords are based on the previous engine ID. For the engine ID, we recommend using the master IP address of the device if the device has multiple routing engines and has the master IP address configured. Alternatively, you can use the MAC address of the management port if the device has only one Routing Engine.

Related Documentation

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
- [Example: SNMPv3 Configuration on page 117](#)

Creating SNMPv3 Users

Supported Platforms [ACX Series](#), [EX4600](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

For each SNMPv3 user, you can specify the username, authentication type, authentication password, privacy type, and privacy password. After a user enters a password, a key based on the engine ID and password is generated and is written to the configuration file. After the generation of the key, the password is deleted from this configuration file.



NOTE: You can configure only one encryption type for each SNMPv3 user.

To create users, include the **user** statement at the **[edit snmp v3 usm local-engine]** hierarchy level:

```
[edit snmp v3 usm local-engine]
user username;
```

username is the name that identifies the SNMPv3 user.

To configure user authentication and encryption, include the following statements at the **[edit snmp v3 usm local-engine user username]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]
authentication-md5 {
  authentication-password authentication-password;
}
authentication-sha {
  authentication-password authentication-password;
}
authentication-none;
privacy-aes128 {
  privacy-password privacy-password;
}
privacy-des {
  privacy-password privacy-password;
}
```

```
privacy-3des {  
  privacy-password privacy-password;  
}  
privacy-none;
```

- Related Documentation**
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
 - [Example: Creating SNMPv3 Users on page 122](#)
 - [Example: SNMPv3 Configuration on page 117](#)

Example: Creating SNMPv3 Users

Define SNMPv3 users:

```
[edit]  
snmp {  
  v3 {  
    usm {  
      local-engine {  
        user user1 {  
          authentication-md5 {  
            authentication-password authentication-password;  
          }  
          privacy-des {  
            privacy-password password;  
          }  
        }  
      }  
      user user2 {  
        authentication-sha {  
          authentication-password authentication-password;  
        }  
        privacy-none;  
      }  
      user user3 {  
        authentication-none;  
        privacy-none;  
      }  
      user user4 {  
        authentication-md5 {  
          authentication-password authentication-password;  
        }  
        privacy-des {  
          privacy-password authentication-password;  
        }  
      }  
      user user5 {  
        authentication-sha {  
          authentication-password authentication-password;  
        }  
        privacy-aes128 {  
          privacy-password authentication-password;  
        }  
      }  
    }  
  }  
}
```

```

    }
  }
}

```

Related Documentation • [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the SNMPv3 Authentication Type

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

By default, in a Junos OS configuration the SNMPv3 authentication type is set to none.

This topic includes the following sections:

- [Configuring MD5 Authentication on page 123](#)
- [Configuring SHA Authentication on page 123](#)
- [Configuring No Authentication on page 124](#)

Configuring MD5 Authentication

To configure the message digest algorithm (MD5) as the authentication type for an SNMPv3 user, include the **authentication-md5** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```

[edit snmp v3 usm local-engine user username]
authentication-md5 {
  authentication-password authentication-password;
}

```

authentication-password is the password used to generate the key used for authentication.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Configuring SHA Authentication

To configure the secure hash algorithm (SHA) as the authentication type for an SNMPv3 user, include the **authentication-sha** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```

[edit snmp v3 usm local-engine user username]
authentication-sha {
  authentication-password authentication-password;
}

```

authentication-password is the password used to generate the key used for authentication.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Configuring No Authentication

To configure no authentication for an SNMPv3 user, include the **authentication-none** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]  
authentication-none;
```

Related Documentation

- [Configuring the SNMPv3 Encryption Type on page 124](#)
- [Defining Access Privileges for an SNMP Group on page 126](#)
- [Configuring the Access Privileges Granted to a Group on page 127](#)
- [Assigning Security Model and Security Name to a Group on page 131](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the SNMPv3 Encryption Type

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

By default, encryption is set to none.



NOTE: Before you configure encryption, you must configure MD5 or SHA authentication.

Before you configure the **privacy-des**, **privacy-3des** and **privacy-aes128** statements, you must install the **jcrypto** package, and either restart the SNMP process or reboot the router.

This topic includes the following sections:

- [Configuring the Advanced Encryption Standard Algorithm on page 124](#)
- [Configuring the Data Encryption Algorithm on page 125](#)
- [Configuring Triple DES on page 125](#)
- [Configuring No Encryption on page 126](#)

Configuring the Advanced Encryption Standard Algorithm

To configure the Advanced Encryption Standard (AES) algorithm for an SNMPv3 user, include the **privacy-aes128** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]  
privacy-aes128 {  
    privacy-password privacy-password;
```

```
}
```

privacy-password is the password used to generate the key used for encryption.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Configuring the Data Encryption Algorithm

To configure the data encryption algorithm (DES) for an SNMPv3 user, include the **privacy-des** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]
privacy-des {
  privacy-password privacy-password;
}
```

privacy-password is the password used to generate the key used for encryption.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Configuring Triple DES

To configure triple DES for an SNMPv3 user, include the **privacy-3des** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]
privacy-3des {
  privacy-password privacy-password;
}
```

privacy-password is the password used to generate the key used for encryption.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Configuring No Encryption

To configure no encryption for an SNMPv3 user, include the **privacy-none** statement at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level:

```
[edit snmp v3 usm local-engine user username]  
privacy-none;
```

Related Documentation

- [Configuring the SNMPv3 Authentication Type on page 123](#)
- [Defining Access Privileges for an SNMP Group on page 126](#)
- [Configuring the Access Privileges Granted to a Group on page 127](#)
- [Assigning Security Model and Security Name to a Group on page 131](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Defining Access Privileges for an SNMP Group

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

The SNMP version 3 (SNMPv3) uses the view-based access control model (VACM), which allows you to configure the access privileges granted to a group. Access is controlled by filtering the MIB objects available for a specific operation through a predefined view. You assign views to determine the objects that are visible for read, write, and notify operations for a particular group, using a particular context, a particular security model (v1, v2c, or usm), and particular security level (authenticated, privacy, or none). For information about how to configure views, see [“Configuring MIB Views” on page 111](#).

You define user access to management information at the **[edit snmp v3 vacm]** hierarchy level. All access control within VACM operates on groups, which are collections of users as defined by USM, or community strings as defined in the SNMPv1 and SNMPv2c security models. The term **security-name** refers to these generic end users. The group to which a specific security name belongs is configured at the **[edit snmp v3 vacm security-to-group]** hierarchy level. That security name can be associated with a group defined at the **[edit snmp v3 vacm security-to-group]** hierarchy level. A group identifies a collection of SNMP users that share the same access policy. You then define the access privileges associated with a group at the **[edit snmp v3 vacm access]** hierarchy level. Access privileges are defined using views. For each group, you can apply different views depending on the SNMP operation; for example, read (**get**, **getNext**, or **getBulk**) write (**set**), notifications, the security level used (authentication, privacy, or none), and the security model (v1, v2c, or usm) used within an SNMP request.

You configure members of a group with the **security-name** statement. For v3 packets using USM, the security name is the same as the username. For SNMPv1 or SNMPv2c packets, the security name is determined based on the community string. Security names are specific to a security model. If you are also configuring VACM access policies for SNMPv1 or SNMPv2c packets, you must assign security names to groups for each security model (SNMPv1 or SNMPv2c) at the **[edit snmp v3 vacm security-to-group]** hierarchy

level. You must also associate a security name with an SNMP community at the **[edit snmp v3 snmp-community *community-index*]** hierarchy level.

To configure the access privileges for an SNMP group, include statements at the **[edit snmp v3 vacm]** hierarchy level:

```
[edit snmp v3 vacm]
access {
  group group-name {
    (default-context-prefix | context-prefix context-prefix){
      security-model (any | usm | v1 | v2c) {
        security-level (authentication | none | privacy) {
          notify-view view-name;
          read-view view-name;
          write-view view-name;
        }
      }
    }
  }
}
security-to-group {
  security-model (usm | v1 | v2c) {
    security-name security-name {
      group group-name;
    }
  }
}
```

Related Documentation

- [Configuring the SNMPv3 Authentication Type on page 123](#)
- [Configuring the Access Privileges Granted to a Group on page 127](#)
- [Assigning Security Model and Security Name to a Group on page 131](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the Access Privileges Granted to a Group

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

This topic includes the following sections:

- [Configuring the Group on page 127](#)
- [Configuring the Security Model on page 128](#)
- [Configuring the Security Level on page 128](#)
- [Associating MIB Views with an SNMP User Group on page 129](#)

Configuring the Group

To configure the access privileges granted to a group, include the **group** statement at the **[edit snmp v3 vacm access]** hierarchy level:

```
[edit snmp v3 vacm access]
group group-name;
```

group-name is a collection of SNMP users that belong to a common SNMP list that defines an access policy. Users belonging to a particular SNMP group inherit all access privileges granted to that group.

Configuring the Security Model

To configure the security model, include the **security-model** statement at the **[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix)]** hierarchy level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
context-prefix)]
security-model (any | usm | v1 | v2c);
```

- **any**—Any security model
- **usm**—SNMPv3 security model
- **v1**—SNMPv1 security model
- **v2c**—SNMPv2c security model

Configuring the Security Level

To configure the access privileges granted to packets with a particular security level, include the **security-level** statement at the **[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix) security-model (any | usm | v1 | v2c)]** hierarchy level:

```
[edit snmp v3 vacm access group group-name default-context-prefix security-model (any
| usm | v1 | v2c)]
security-level (authentication | none | privacy);
```

- **none**—Provides no authentication and no encryption.
- **authentication**—Provides authentication but no encryption.
- **privacy**—Provides authentication and encryption.



NOTE: Access privileges are granted to all packets with a security level equal to or greater than that configured. If you are configuring the SNMPv1 or SNMPv2c security model, use **none** as your security level. If you are configuring the SNMPv3 security model (USM), use the **authentication**, **none**, or **privacy** security level.

- See Also**
- [Configuring the SNMPv3 Authentication Type on page 123](#)
 - [Defining Access Privileges for an SNMP Group on page 126](#)
 - [Assigning Security Model and Security Name to a Group on page 131](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Associating MIB Views with an SNMP User Group

MIB views define access privileges for members of a group. Separate views can be applied for each SNMP operation (read, write, and notify) within each security model (usm, v1, and v2c) and each security level (authentication, none, and privacy) supported by SNMP.

To associate MIB views with an SNMP user group, include the following statements at the `[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]` hierarchy level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
  context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none
  | privacy)]
  notify-view view-name;
  read-view view-name;
  write-view view-name;
```



NOTE: You must associate at least one view (notify, read, or write) at the `[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]` hierarchy level.

You must configure the MIB view at the `[edit snmp view view-name]` hierarchy level. For information about how to configure MIB views, see [“Configuring MIB Views” on page 111](#).

This section describes the following topics related to this configuration:

- [Configuring the Notify View on page 129](#)
- [Configuring the Read View on page 130](#)
- [Configuring the Write View on page 130](#)

Configuring the Notify View

To associate notify access with an SNMP user group, include the **notify-view** statement at the `[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]` hierarchy level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
  context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none
  | privacy)]
  notify-view view-name;
```

view-name specifies the notify access, which is a list of notifications that can be sent to each user in an SNMP group. A view name cannot exceed 32 characters.

Configuring the Read View

To associate a read view with an SNMP group, include the **read-view** statement at the **[edit snmp v3 vacm access group *group-name* (default-context-prefix | context-prefix *context-prefix*) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]** hierarchy level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
  context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none
    | privacy)]
  read-view view-name;
```

view-name specifies read access for an SNMP user group. A view name cannot exceed 32 characters.

Configuring the Write View

To associate a write view with an SNMP user group, include the **write-view** statement at the **[edit snmp v3 vacm access group *group-name* (default-context-prefix | context-prefix *context-prefix*) security-model (any | usm | v1 | v2c) security-level (authentication | none | privacy)]** hierarchy level:

```
[edit snmp v3 vacm access group group-name (default-context-prefix | context-prefix
  context-prefix) security-model (any | usm | v1 | v2c) security-level (authentication | none
    | privacy)]
  write-view view-name;
```

view-name specifies write access for an SNMP user group. A view name cannot exceed 32 characters.

Related Documentation

- [Configuring the SNMPv3 Authentication Type on page 123](#)
- [Defining Access Privileges for an SNMP Group on page 126](#)
- [Assigning Security Model and Security Name to a Group on page 131](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
- [Example: Configuring the Access Privileges Granted to a Group on page 130](#)

Example: Configuring the Access Privileges Granted to a Group

Supported Platforms **ACX Series, M Series, MX Series, PTX Series, T Series**

Define access privileges:

```
[edit snmp v3 vacm]
access {
  group group1 {
    default-context-prefix {
      security-model usm {          #Define an SNMPv3 security model
        security-level privacy {
          notify-view nv1;
          read-view rv1;
          write-view wv1;
        }
      }
    }
  }
}
```

```

    }
  }
  context-prefix lr1/ri1 { # routing instance ri1 in logical system lr1
    security-model usm {
      security-level privacy {
        notify-view nv1;
        read-view rv1;
        write-view wv1;
      }
    }
  }
}
group group2 {
  default-context-prefix {
    security-model usm {      #Define an SNMPv3 security model
      security-level authentication {
        read-view rv2;
        write-view wv2;
      }
    }
  }
}
group group3 {
  default-context-prefix {
    security-model v1 {      #Define an SNMPv3 security model
      security-level none {
        read-view rv3;
        write-view wv3;
      }
    }
  }
}
}

```

- Related Documentation**
- [Configuring the Access Privileges Granted to a Group on page 127](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Assigning Security Model and Security Name to a Group

Supported Platforms ACX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX

To assign security names to groups, include the following statements at the **[edit snmp v3 vacm security-to-group]** hierarchy level:

```

[edit snmp v3 vacm security-to-group]
security-model (usm | v1 | v2c) {
  security-name security-name {
    group group-name;
  }
}

```

This topic includes the following sections:

- [Configuring the Security Model on page 132](#)
- [Assigning Security Names to Groups on page 132](#)
- [Configuring the Group on page 132](#)

Configuring the Security Model

To configure the security model, include the **security-model** statement at the **[edit snmp v3 vacm security-to-group]** hierarchy level:

```
[edit snmp v3 vacm security-to-group]  
security-model (usm | v1 | v2c);
```

- **usm**—SNMPv3 security model
- **v1**—SNMPv1 security model
- **v2c**—SNMPv2 security model

Assigning Security Names to Groups

To associate a security name with an SNMPv3 user, or a v1 or v2 community string, include the **security-name** statement at the **[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]** hierarchy level:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]  
security-name security-name;
```

For SNMPv3, the **security-name** is the username configured at the **[edit snmp v3 usm local-engine user username]** hierarchy level. For SNMPv1 and SNMPv2c, the security name is the community string configured at the **[edit snmp v3 snmp-community community-index]** hierarchy level. For information about configuring usernames, see [“Creating SNMPv3 Users” on page 121](#). For information about configuring a community string, see [“Configuring the SNMPv3 Community” on page 151](#).



NOTE: The USM security name is separate from the SNMPv1 and SNMPv2c security name. If you support SNMPv1 and SNMPv2c in addition to SNMPv3, you must configure separate security names within the security-to-group configuration at the **[edit snmp v3 vacm access]** hierarchy level.

Configuring the Group

After you have created SNMPv3 users, or v1 or v2 security names, you associate them with a group. A group is a set of security names belonging to a particular security model. A group defines the access rights for all users belonging to it. Access rights define what SNMP objects can be read, written to, or created. A group also defines what notifications a user is allowed to receive.

If you already have a group that is configured with all of the view and access permissions that you want to give a user, you can add the user to that group. If you want to give a user

view and access permissions that no other groups have, or if you do not have any groups configured, create a group and add the user to it.

To configure the access privileges granted to a group, include the **group** statement at the **[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name security-name]** hierarchy level:

```
[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name security-name]
group group-name;
```

group-name identifies a collection of SNMP security names that share the same access policy. For more information about groups, see “Defining Access Privileges for an SNMP Group” on page 126.

- See Also**
- [Configuring the SNMPv3 Authentication Type on page 123](#)
 - [Defining Access Privileges for an SNMP Group on page 126](#)
 - [Configuring the Access Privileges Granted to a Group on page 127](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
 - [Example: Security Group Configuration on page 133](#)

Example: Security Group Configuration

Supported Platforms [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Assign security names to groups:

```
vacm {
  security-to-group {
    security-model usm {
      security-name user1 {
        group group1;
      }
      security-name user2 {
        group group2;
      }
      security-name user3 {
        group group3;
      }
    }
  }
}
```

- Related Documentation**
- [Assigning Security Model and Security Name to a Group on page 131](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring SNMPv3 Traps on a Device Running Junos OS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

In SNMPv3, you create traps and informs by configuring the **notify**, **target-address**, and **target-parameters** parameters. Traps are unconfirmed notifications, whereas informs are confirmed notifications. This section describes how to configure SNMP traps. For information about configuring SNMP informs, see [“Configuring SNMP Informs” on page 143](#).

The target address defines a management application's address and parameters to be used in sending notifications. Target parameters define the message processing and security parameters that are used in sending notifications to a particular management target. SNMPv3 also lets you define SNMPv1 and SNMPv2c traps.



NOTE: When you configure SNMP traps, make sure your configured access privileges allow the traps to be sent. Access privileges are configured at the `[edit snmp v3 vacm access]` and `[edit snmp v3 vacm security-to-group]` hierarchy levels.

To configure SNMP traps, include the following statements at the `[edit snmp v3]` hierarchy level:

```
[edit snmp v3]
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter name {
    oid object-identifier (include | exclude);
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    routing-instance instance;
    tag-list tag-list;
    target-parameters target-parameters-name;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | v3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
}
```

Related Documentation

- [Configuring the SNMPv3 Trap Notification on page 135](#)
- [Configuring the Trap Notification Filter on page 136](#)
- [Configuring the Trap Target Address on page 137](#)
- [Defining and Configuring the Trap Target Parameters on page 140](#)

- [Configuring SNMP Informs on page 143](#)
- [Configuring the Remote Engine and Remote User on page 144](#)
- [Configuring the Inform Notification Type and Target Address on page 148](#)

Configuring the SNMPv3 Trap Notification

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

The **notify** statement specifies the type of notification (trap) and contains a single tag. The tag defines a set of target addresses to receive a trap. The tag list contains one or more tags and is configured at the **[edit snmp v3 target-address *target-address-name*]** hierarchy level. If the tag list contains this tag, Junos OS sends a notification to all the target addresses associated with this tag.

To configure the trap notifications, include the **notify** statement at the **[edit snmp v3]** hierarchy level:

```
[edit snmp v3]
notify name {
  tag tag-name;
  type trap;
}
```

name is the name assigned to the notification.

tag-name defines the target addresses to which this notification is sent. This notification is sent to all the target-addresses that have this tag in their tag list. The **tag-name** is not included in the notification.

trap is the type of notification.



NOTE: Each notify entry name must be unique.

Junos OS supports two types of notification: **trap** and **inform**.

For information about how to configure the tag list, see [“Configuring the Trap Target Address” on page 139](#).

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring the Trap Notification Filter on page 136](#)
- [Configuring the Trap Target Address on page 137](#)
- [Defining and Configuring the Trap Target Parameters on page 140](#)
- [Configuring SNMP Informs on page 143](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Example: Configuring SNMPv3 Trap Notification

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Specify three sets of destinations to send traps:

```
[edit snmp v3]
  notify n1 {
    tag router1;
    type trap;
  }
  notify n2 {
    tag router2;
    type trap;
  }
  notify n3 {
    tag router3;
    type trap;
  }
```

- Related Documentation**
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the Trap Notification Filter

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

SNMPv3 uses the notify filter to define which traps (or which objects from which traps) are sent to the network management system (NMS). The trap notification filter limits the type of traps that are sent to the NMS.

Each object identifier represents a subtree of the MIB object hierarchy. The subtree can be represented either by a sequence of dotted integers (such as **1.3.6.1.2.1.2**) or by its subtree name (such as **interfaces**). You can also use the wildcard character asterisk (*) in the object identifier (OID) to specify object identifiers that match a particular pattern.

To configure the trap notifications filter, include the **notify-filter** statement at the **[edit snmp v3]** hierarchy level:

```
[edit snmp v3]
  notify-filter profile-name;
```

profile-name is the name assigned to the notify filter.

By default, the OID is set to **include**. To define access to traps (or objects from traps), include the **oid** statement at the **[edit snmp v3 notify-filter *profile-name*]** hierarchy level:

```
[edit snmp v3 notify-filter profile-name]
  oid oid (include | exclude);
```

oid is the object identifier. All MIB objects represented by this statement have the specified OID as a prefix. It can be specified either by a sequence of dotted integers or by a subtree name.

- **include**—Include the subtree of MIB objects represented by the specified OID.
- **exclude**—Exclude the subtree of MIB objects represented by the specified OID.

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring the SNMPv3 Trap Notification on page 135](#)
- [Configuring the Trap Target Address on page 137](#)
- [Defining and Configuring the Trap Target Parameters on page 140](#)
- [Configuring SNMP Informs on page 143](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the Trap Target Address

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

The target address defines a management application's address and parameters that are used in sending notifications. It can also identify management stations that are allowed to use specific community strings. When you receive a packet with a recognized community string and a tag is associated with it, Junos OS looks up all the target addresses with this tag and verifies that the source address of this packet matches one of the configured target addresses.



NOTE: You must configure the address mask when you configure the SNMP community.

To specify where you want the traps to be sent and define what SNMPv1 and SNMPv2cc packets are allowed, include the **target-address** statement at the **[edit snmp v3]** hierarchy level:

```
[edit snmp v3]
target-address target-address-name;
```

target-address-name is the string that identifies the target address.

To configure the target address properties, include the following statements at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]
address address;
address-mask address-mask;
logical-system logical-system;
port port-number;
routing-instance instance;
tag-list tag-list;
target-parameters target-parameters-name;
```

This section includes the following topics:

- [Configuring the Address on page 138](#)
- [Configuring the Address Mask on page 138](#)
- [Configuring the Port on page 138](#)
- [Configuring the Routing Instance on page 138](#)
- [Configuring the Trap Target Address on page 139](#)
- [Applying Target Parameters on page 139](#)

Configuring the Address

To configure the address, include the **address** statement at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]  
address address;
```

address is the SNMP target address.

Configuring the Address Mask

The address mask specifies a set of addresses that are allowed to use a community string and verifies the source addresses for a group of target addresses.

To configure the address mask, include the **address-mask** statement at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]  
address-mask address-mask;
```

address-mask combined with the address defines a range of addresses. For information about how to configure the community string, see [“Configuring the SNMPv3 Community” on page 151](#).

Configuring the Port

By default, the UDP port is set to 162. To configure a different port number, include the **port** statement at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]  
port port-number;
```

port-number is the SNMP target port number.

Configuring the Routing Instance

Traps are sent over the default routing instance. To configure the routing instance for sending traps, include the **routing-instance** statement at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]  
routing-instance instance;
```

instance is the name of the routing instance. To configure a routing instance within a logical system, specify the logical system name followed by the routing instance name.

Use a slash (/) to separate the two names (for example, **test-lr/test-ri**). To configure the default routing instance on a logical system, specify the logical system name followed by **default** (for example, **test-lr/default**).

Configuring the Trap Target Address

Each **target-address** statement can have one or more tags configured in its tag list. Each tag can appear in more than one tag list. When a significant event occurs on the network device, the tag list identifies the targets to which a notification is sent.

To configure the tag list, include the **tag-list** statement at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]
tag-list "tag-list";
```

tag-list specifies one or more tags as a space-separated list enclosed within double quotes.

For an example of tag list configuration, see [“Example: Configuring the Tag List” on page 140](#).

For information about how to specify a tag at the **[edit snmp v3 notify notify-name]** hierarchy level, see [“Configuring the SNMPv3 Trap Notification” on page 135](#).



NOTE: When you configure SNMP traps, make sure your configured access privileges allow the traps to be sent. Configure access privileges at the **[edit snmp v3 vacm access]** hierarchy level.

Applying Target Parameters

The **target-parameters** statement at the **[edit snmp v3]** hierarchy level applies the target parameters configured at the **[edit snmp v3 target-parameters target-parameters-name]** hierarchy level.

To reference configured target parameters, include the **target-parameters** statement at the **[edit snmp v3 target-address target-address-name]** hierarchy level:

```
[edit snmp v3 target-address target-address-name]
target-parameters target-parameters-name;
```

target-parameters-name is the name associated with the message processing and security parameters that are used in sending notifications to a particular management target.

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring the SNMPv3 Trap Notification on page 135](#)
- [Configuring the Trap Notification Filter on page 136](#)
- [Defining and Configuring the Trap Target Parameters on page 140](#)
- [Configuring SNMP Informs on page 143](#)

- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
- [Example: Configuring the Tag List on page 140](#)

Example: Configuring the Tag List

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

In the following example, two tag entries (**router1** and **router2**) are defined at the **[edit snmp v3 notify *notify-name*]** hierarchy level. When an event triggers a notification, Junos OS sends a trap to all target addresses that have **router1** or **router2** configured in their target-address tag list. This results in the first two targets getting one trap each, and the third target getting two traps.

```
[edit snmp v3]
notify n1 {
  tag router1; # Identifies a set of target addresses
  type trap; # Defines the type of notification
}
notify n2 {
  tag router2;
  type trap;
}
target-address ta1 {
  address 10.1.1.1;
  address-mask 255.255.255.0;
  port 162;
  tag-list router1;
  target-parameters tp1;
}
target-address ta2 {
  address 10.1.1.2;
  address-mask 255.255.255.0;
  port 162;
  tag-list router2;
  target-parameters tp2;
}
target-address ta3 {
  address 10.1.1.3;
  address-mask 255.255.255.0;
  port 162;
  tag-list "router1 router2"; #Define multiple tags in the target address tag list
  target-parameters tp3;
}
```

- Related Documentation**
- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
 - [Configuring the Trap Target Address on page 137](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Defining and Configuring the Trap Target Parameters

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

Target parameters define the message processing and security parameters that are used in sending notifications to a particular management target.

To define a set of target parameters, include the **target-parameters** statement at the **[edit snmp v3]** hierarchy level:

```
[edit snmp v3]
target-parameters target-parameters-name;
```

target-parameters-name is the name assigned to the target parameters.

To configure target parameter properties, include the following statements at the **[edit snmp v3 target-parameters target-parameter-name]** hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name]
notify-filter profile-name;
parameters {
  message-processing-model (v1 | v2c | V3);
  security-level (authentication | none | privacy);
  security-model (usm | v1 | v2c);
  security-name security-name;
}
```

This topic includes the following sections:

- [Applying the Trap Notification Filter on page 141](#)
- [Configuring the Target Parameters on page 141](#)

Applying the Trap Notification Filter

To apply the trap notification filter, include the **notify-filter** statement at the **[edit snmp v3 target-parameters target-parameter-name]** hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name]
notify-filter profile-name;
```

profile-name is the name of a configured notify filter. For information about configuring notify filters, see [“Configuring the Trap Notification Filter” on page 136](#).

Configuring the Target Parameters

To configure target parameter properties, include the following statements at the **[edit snmp v3 target-parameters target-parameter-name parameters]** hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name parameters]
message-processing-model (v1 | v2c | v3);
security-level (authentication | none | privacy);
security-model (usm | v1 | v2c);
security-name security-name;
```

This section includes the following topics:

- [Configuring the Message Processing Model on page 142](#)
- [Configuring the Security Model on page 142](#)
- [Configuring the Security Level on page 142](#)
- [Configuring the Security Name on page 143](#)

Configuring the Message Processing Model

The message processing model defines which version of SNMP to use when generating SNMP notifications. To configure the message processing model, include the **message-processing-model** statement at the **[edit snmp v3 target-parameters target-parameter-name parameters]** hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name parameters]
  message-processing-model (v1 | v2c | v3);
```

- **v1**—SNMPv1 message processing model
- **v2c**—SNMPv2c message processing model
- **v3**—SNMPv3 message processing model

Configuring the Security Model

To define the security model to use when generating SNMP notifications, include the **security-model** statement at the **[edit snmp v3 target-parameters target-parameter-name parameters]** hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name parameters]
  security-model (usm | v1 | v2c);
```

- **usm**—SNMPv3 security model
- **v1**—SNMPv1 security model
- **v2c**—SNMPv2c security model

Configuring the Security Level

The **security-level** statement specifies whether the trap is authenticated and encrypted before it is sent.

To configure the security level to use when generating SNMP notifications, include the **security-level** statement at the **[edit snmp v3 target-parameters target-parameter-name parameters]** hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name parameters]
  security-level (authentication | none | privacy);
```

- **authentication**—Provides authentication but no encryption.
- **none**—No security. Provides no authentication and no encryption.
- **privacy**—Provides authentication and encryption.



NOTE: If you are configuring the SNMPv1 or SNMPv2c security model, use **none** as your security level. If you are configuring the SNMPv3 (USM) security model, use the **authentication** or **privacy** security level.

Configuring the Security Name

To configure the security name to use when generating SNMP notifications, include the **security-name** statement at the `[edit snmp v3 target-parameters target-parameter-name parameters]` hierarchy level:

```
[edit snmp v3 target-parameters target-parameter-name parameters]
  security-name security-name;
```

If the USM security model is used, the **security-name** identifies the user that is used when the notification is generated. If the v1 or v2c security models are used, **security-name** identifies the SNMP community used when the notification is generated.



NOTE: The access privileges for the group associated with a security name must allow this notification to be sent.

If you are using the v1 or v2 security models, the security name at the `[edit snmp v3 vacm security-to-group]` hierarchy level must match the security name at the `[edit snmp v3 snmp-community community-index]` hierarchy level.

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring the SNMPv3 Trap Notification on page 135](#)
- [Configuring the Trap Notification Filter on page 136](#)
- [Configuring the Trap Target Address on page 137](#)
- [Configuring SNMP Informs on page 143](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring SNMP Informs

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

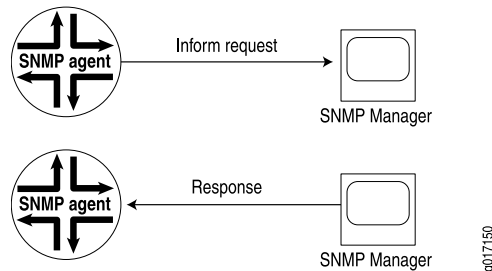
Junos OS supports two types of notifications: traps and informs. With traps, the receiver does not send any acknowledgment when it receives a trap. Therefore, the sender cannot determine if the trap was received. A trap may be lost because a problem occurred during transmission. To increase reliability, an inform is similar to a trap except that the inform is stored and retransmitted at regular intervals until one of these conditions occurs:

- The receiver (target) of the inform returns an acknowledgment to the SNMP agent.
- A specified number of unsuccessful retransmissions have been attempted and the agent discards the inform message.

If the sender never receives a response, the inform can be sent again. Thus, informs are more likely to reach their intended destination than traps are. Informs use the same communications channel as traps (same socket and port) but have different protocol data unit (PDU) types.

Informs are more reliable than traps, but they consume more network, router, and switch resources (see [Figure 1 on page 144](#)). Unlike a trap, an inform is held in memory until a response is received or the timeout is reached. Also, traps are sent only once, whereas an inform may be retried several times. Use informs when it is important that the SNMP manager receive all notifications. However, if you are more concerned about network traffic, or router and switch memory, use traps.

Figure 1: Inform Request and Response



For information about configuring SNMP traps, see [“Configuring SNMPv3 Traps on a Device Running Junos OS” on page 133](#).

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring the Remote Engine and Remote User on page 144](#)
- [Configuring the Inform Notification Type and Target Address on page 148](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the Remote Engine and Remote User

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

To send inform messages to an SNMPv3 user on a remote device, you must first specify the engine identifier for the SNMP agent on the remote device where the user resides. The remote engine ID is used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host. When sending an inform message, the agent uses the credentials of the user configured on the remote engine (inform target).

To configure a remote engine and remote user to receive and respond to SNMP informs, include the following statements at the **[edit snmp v3]** hierarchy level:

```

[edit snmp v3]
usm {
  remote-engine engine-id {
    user username {
      authentication-md5 {
        authentication-key key;
      }
      authentication-none;
      authentication-sha {
        authentication-key key;
      }
    }
  }
}
  
```

```

    privacy-3des {
        privacy-key key;
    }
    privacy-aes128 {
        privacy-key key;
    }
    privacy-des {
        privacy-key key;
    }
    privacy-none;
}
}
}

```

For informs, **remote-engine *engine-id*** is the identifier for the SNMP agent on the remote device where the user resides.

For informs, **user *username*** is the user on a remote SNMP engine who receives the informs.

Informs generated can be **unauthenticated**, **authenticated**, or **authenticated_and_encrypted**, depending on the security level of the SNMPv3 user configured on the remote engine (the inform receiver). The authentication key is used for generating message authentication code (MAC). The privacy key is used to encrypt the inform PDU part of the message.

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring SNMP Informs on page 143](#)
- [Configuring the Inform Notification Type and Target Address on page 148](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
- [Example: Configuring the Remote Engine ID and Remote User on page 145](#)

Example: Configuring the Remote Engine ID and Remote User

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

This example shows how to configure a remote engine and remote user so you can receive and respond to SNMP inform notifications. Inform notifications can be authenticated and encrypted. They are also more reliable than traps, another type of notification that Junos OS supports. Unlike traps, inform notifications are stored and retransmitted at regular intervals until one of these conditions occurs:

- The target of the inform notification returns an acknowledgment to the SNMP agent.
- A specified number of unsuccessful retransmissions have been attempted.
- [Requirements on page 146](#)
- [Overview on page 146](#)
- [Configuration on page 146](#)
- [Verification on page 148](#)

Requirements

No special configuration beyond device initialization is required before configuring this example.

This feature requires the use of plain-text passwords valid for SNMPv3. SNMPv3 has the following special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Although quotation marks are not always required to enclose passwords, it is best to use them. You need quotation marks if the password contains any spaces or possibly in the case of certain special characters or punctuation.

Overview

Inform notifications are supported in SNMPv3 to increase reliability. For example, an SNMP agent receiving an inform notification acknowledges the receipt.

For inform notifications, the remote engine ID identifies the SNMP agent on the remote device where the user resides, and the username identifies the user on a remote SNMP engine who receives the inform notifications.

Consider a scenario in which you have the values in [Table 12 on page 146](#) to use in configuring the remote engine ID and remote user in this example.

Table 12: Values to Use in Example

Name of Variable	Value
username	u10
remote engine ID	800007E5804089071BC6D10A41
authentication type	authentication-md5
authentication password	qol67R%?
encryption type	privacy-des
privacy password	m*72Jl9v

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands and paste them into a text file, remove any line breaks and change any details necessary to match your network configuration, copy and paste these commands into the CLI at the **[edit snmp v3]** hierarchy level, and then enter **commit** from configuration mode.

```

set usm remote-engine 800007E5804089071BC6D10A41 user u10 authentication-md5
authentication-key "qol67R%?"
set usm remote-engine 800007E5804089071BC6D10A41 user u10 privacy-des privacy-key
"m*72Jl9v"

```

Configuring the Remote Engine and Remote User

Step-by-Step Procedure The following example requires that you navigate to various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *Junos OS CLI User Guide*.

To configure the remote engine ID and remote user:

1. Configure the remote engine ID, username, and authentication type and password.

```

[edit snmp v3]
user@host# set usm remote-engine 800007E5804089071BC6D10A41 user u10
authentication-md5 authentication-key "qol67R%?"

```

2. Configure the encryption type and privacy password.

You can configure only one encryption type per SNMPv3 user.

```

[edit snmp v3]
user@host# set usm remote-engine 800007E5804089071BC6D10A41 user u10
privacy-des privacy-key "m*72Jl9v"

```

Results

In configuration mode, confirm your configuration by entering the **show** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

[edit snmp v3]
user@ host# show
usm {
  remote-engine 800007E5804089071BC6D10A41 {
    user u10 {
      authentication-md5 {
        authentication-key "$9$Tz/teK8NdsLXk.f5n6p0ORev"; ## SECRET-DATA
      }
      privacy-des {
        privacy-key "$9$/gyNCu1KvWdwYMWw2gJHkRhcrWx"; ## SECRET-DATA
      }
    }
  }
}

```

After you have confirmed that the configuration is correct, enter **commit** from configuration mode.

Verification

Verifying the Configuration of the Remote Engine ID and Username

Purpose Verify the status of the engine ID and user information.

Action Display information about the SNMPv3 engine ID and user.

```
user@host> show snmp v3
Local engine ID: 80 00 0a 4c 01 0a ff 03 e3
Engine boots:      3
Engine time:       769187 seconds
Max msg size:      65507 bytes

Engine ID: 80 00 07 e5 80 40 89 07 1b c6 d1 0a 41
  User                               Auth/Priv  Storage    Status
  u10                                md5/des   nonvolatile active
```

Meaning The output displays the following information:

- Local engine ID and detail about the engine
- Remote engine ID (labeled **Engine ID**)
- Username
- Authentication type and encryption (privacy) type that is configured for the user
- Type of storage for the username, either nonvolatile (configuration saved) or volatile (not saved)
- Status of the new user; only users with an active status can use SNMPv3

- Related Documentation**
- [show snmp v3 on page 950](#)
 - [Configuring the SNMPv3 Encryption Type on page 124](#)
 - [Configuring the SNMPv3 Authentication Type on page 123](#)
 - [Configuring SNMP Informs on page 143](#)
 - [Configuring the Remote Engine and Remote User on page 144](#)

Configuring the Inform Notification Type and Target Address

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

To configure the inform notification type and target information, include the following statements at the **[edit snmp v3]** hierarchy level:

```
[edit snmp v3]
  notify name {
    tag tag-name;
```

```

    type (trap | inform);
}
target-address target-address-name {
    address address;
    address-mask address-mask;
    logical-system logical-system;
    port port-number;
    retry-count number;
    routing-instance instance;
    tag-list tag-list;
    target-parameters target-parameters-name;
    timeout seconds;
}
target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
        message-processing-model (v1 | v2c | v3);
        security-level (authentication | none | privacy);
        security-model (usm | v1 | v2c);
        security-name security-name;
    }
}

```

notify *name* is the name assigned to the notification. Each notify entry name must be unique.

tag *tag-name* defines the target addresses that are sent this notification. The notification is sent to all target addresses that have this tag in their tag list. The **tag-name** is not included in the notification. For information about how to configure the tag list, see [“Configuring the Trap Target Address” on page 139](#).

type *inform* is the type of notification.

target-address *target-address-name* identifies the target address. The target address defines a management application’s address and parameters that are used to respond to informs.

timeout *seconds* is the number of seconds to wait for an acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted. The default timeout is 15 seconds.

retry-count *number* is the maximum number of times an inform is transmitted if no acknowledgment is received. The default is 3. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded.

message-processing-model defines which version of SNMP to use when SNMP notifications are generated. Informs require a **v3** message processing model.

security-model defines the security model to use when SNMP notifications are generated. Informs require a **usm** security model.

security-model defines the security model to use when SNMP notifications are generated. Informs require a **usm** security model.

security-level specifies whether the inform is authenticated and encrypted before it is sent. For the **usm** security model, the security level must be one of the following:

- **authentication**—Provides authentication but no encryption.
- **privacy**—Provides authentication and encryption.

security-name identifies the username that is used when generating the inform.

Related Documentation

- [Configuring SNMPv3 Traps on a Device Running Junos OS on page 133](#)
- [Configuring SNMP Informs on page 143](#)
- [Configuring the Remote Engine and Remote User on page 144](#)
- [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
- [Example: Configuring the Inform Notification Type and Target Address on page 150](#)

Example: Configuring the Inform Notification Type and Target Address

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

In the following example, target **172.17.20.184** is configured to respond to informs. The inform timeout is **30** seconds and the maximum retransmit count is **3**. The inform is sent to all targets in the **tl1** list. The security model for the remote user is **usm** and the remote engine username is **u10**.

```
[edit snmp v3]
notify n1 {
  type inform;
  tag tl1;
}
notify-filter nf1 {
  oid .1.3 include;
}
target-address ta1 {
  address 172.17.20.184;
  retry-count 3;
  tag-list tl1;
  address-mask 255.255.255.0;
  target-parameters tp1;
  timeout 30;
}
target-parameters tp1 {
  parameters {
    message-processing-model v3;
    security-model usm;
    security-level privacy;
    security-name u10;
  }
  notify-filter nf1;
}
```


- Related Documentation**
- [Configuring the Inform Notification Type and Target Address on page 148](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

Configuring the SNMPv3 Community

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

The SNMP community defines the relationship between an SNMP server system and the client systems. This statement is optional.

To configure the SNMP community, include the **snmp-community** statement at the **[edit snmp v3]** hierarchy level:

```
[edit snmp v3]
snmp-community community-index;
```

community-index is the index for the SNMP community.

To configure the SNMP community properties, include the following statements at the **[edit snmp v3 snmp-community community-index]** hierarchy level:

```
[edit snmp v3 snmp-community community-index]
community-name community-name;
context context-name;
security-name security-name;
tag tag-name;
```

This section includes the following topics:

- [Configuring the Community Name on page 151](#)
- [Configuring the Context on page 152](#)
- [Configuring the Security Names on page 152](#)
- [Configuring the Tag on page 152](#)

Configuring the Community Name

The community name defines the SNMP community. The SNMP community authorizes SNMPv1 or SNMPv2c clients. The access privileges associated with the configured security name define which MIB objects are available and the operations (read, write, or notify) allowed on those objects.

To configure the SNMP community name, include the **community-name** statement at the **[edit snmp v3 snmp-community community-index]** hierarchy level:

```
[edit snmp v3 snmp-community community-index]
community-name community-name;
```

community-name is the community string for an SNMPv1 or SNMPv2c community.

If unconfigured, it is the same as the community index.

If the community name contains spaces, enclose it in quotation marks (" ").



NOTE: Community names must be unique. You cannot configure the same community name at the [edit snmp community] and [edit snmp v3 snmp-community *community-index*] hierarchy levels. The configured community name at the [edit snmp v3 snmp-community *community-index*] hierarchy level is encrypted. You cannot view the community name after you have configured it and committed your changes. In the command-line interface (CLI), the community name is concealed.

Configuring the Context

An SNMP context defines a collection of management information that is accessible to an SNMP entity. Typically, an SNMP entity has access to multiple contexts. A context can be a physical or logical system, a collection of multiple systems, or even a subset of a system. Each context in a management domain has a unique identifier.

To configure an SNMP context, include the **context context-name** statement at the [edit snmp v3 snmp-community *community-index*] hierarchy level:

```
[edit snmp v3 snmp-community community-index]  
context context-name;
```



NOTE: To query a routing instance or a logical system,

Configuring the Security Names

To assign a community string to a security name, include the **security-name** statement at the [edit snmp v3 snmp-community *community-index*] hierarchy level:

```
[edit snmp v3 snmp-community community-index]  
security-name security-name;
```

security-name is used when access control is set up. The **security-to-group** configuration at the [edit snmp v3 vacm] hierarchy level identifies the group.



NOTE: This security name must match the security name configured at the [edit snmp v3 target-parameters *target-parameters-name* parameters] hierarchy level when you configure traps.

Configuring the Tag

To configure the tag, include the **tag** statement at the [edit snmp v3 snmp-community *community-index*] hierarchy level:

```
[edit snmp v3 snmp-community community-index]  
tag tag-name;
```

tag-name identifies the address of managers that are allowed to use a community string.

- Related Documentation**
- [Creating SNMPv3 Users on page 121](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)
 - [Example: Configuring an SNMPv3 Community on page 153](#)

Example: Configuring an SNMPv3 Community

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Define an SNMP community:

```
[edit snmp v3]
snmp-community index1 {
  community-name "$9$JOZl.QF/AtOz3"; # SECRET-DATA
  security-name john;
  tag router1; # Identifies managers that are allowed to use
               # a community string
  target-address ta1 {
    address 10.1.1.1;
    address-mask 255.255.255.0; # Defines the range of addresses
    port 162;
    tag-list router1;
    target-parameters tp1; # Applies configured target parameters
  }
}
```

- Related Documentation**
- [Configuring the SNMPv3 Community on page 151](#)
 - [Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116](#)

CHAPTER 8

Configuring SNMP for Routing Instances

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [SNMP MIBs Supported for Routing Instances on page 156](#)
- [Support Classes for MIB Objects on page 166](#)
- [SNMP Traps Supported for Routing Instances on page 167](#)
- [Identifying a Routing Instance on page 168](#)
- [Enabling SNMP Access over Routing Instances on page 169](#)
- [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)
- [Example: Configuring Interface Settings for a Routing Instance on page 170](#)
- [Configuring Access Lists for SNMP Access over Routing Instances on page 172](#)

Understanding SNMP Support for Routing Instances

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Junos OS enables SNMP managers for all routing instances to request and manage SNMP data related to the corresponding routing instances and logical system networks.

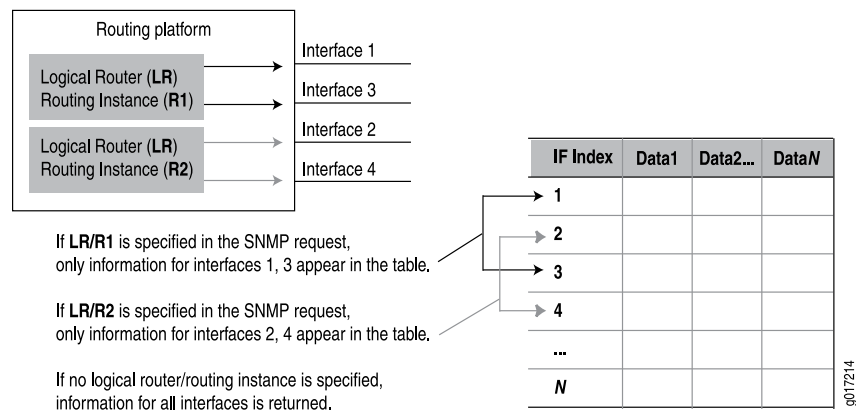
In Junos OS:

- Clients from routing instances other than the default can access MIB objects and perform SNMP operations only on the logical system networks to which they belong.
- Clients from the default routing instance can access information related to all routing instances and logical system networks.

Before Junos OS Release 8.4, only the SNMP manager in the default routing instance (**inet.0**) had access to the MIB objects

With the increase in virtual private network (VPN) service offerings, this feature is useful particularly for service providers who need to obtain SNMP data for specific routing instances (see [Figure 2 on page 156](#)). Service providers can use this information for their own management needs or export the data for use by their customers.

Figure 2: SNMP Data for Routing Instances



If no routing instance is specified in the request, the SNMP agent operates as before:

- For nonrouting table objects, all instances are exposed.
- For routing table objects, only those associated with the default routing instance are exposed.



NOTE: The actual protocol data units (PDUs) are still exchanged over the default (inet.0) routing instance, but the data contents returned are dictated by the routing instance specified in the request PDUs.

Related Documentation

- [Support Classes for MIB Objects on page 166](#)
- [SNMP Traps Supported for Routing Instances on page 167](#)
- [Identifying a Routing Instance on page 168](#)
- [Enabling SNMP Access over Routing Instances on page 169](#)
- [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)
- [Configuring Access Lists for SNMP Access over Routing Instances on page 172](#)

SNMP MIBs Supported for Routing Instances

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

[Table 13 on page 156](#) shows enterprise-specific MIB objects supported by Junos OS and provides notes detailing how they are handled when a routing instance is specified in an SNMP request. An en dash (–) indicates that the item is not applicable.

Table 13: MIB Support for Routing Instances (Juniper Networks MIBs)

Object	Support Class	Description/Notes
jnxProducts(1)	–	Product Object IDs

Table 13: MIB Support for Routing Instances (Juniper Networks MIBs) (*continued*)

Object	Support Class	Description/Notes
jnxServices(2)	–	Services
jnxMibs(3) jnxBoxAnatomy(1)	Class 3	Objects are exposed only for the default logical system.
mpls(2)	Class 2	All instances within a logical system are exposed. Data will not be segregated down to the routing instance level.
ifJnx(3)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxAlarms(4)	Class 3	Objects are exposed only for the default logical system.
jnxFirewalls(5)	Class 4	Data is not segregated by routing instance. All instances are exposed.
jnxDCUs(6)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxPingMIB(7)	Class 3	Objects are exposed only for the default logical system.
jnxTraceRouteMIB(8)	Class 3	Objects are exposed only for the default logical system.
jnxATM(10)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxIpv6(11)	Class 4	Data is not segregated by routing instance. All instances are exposed.
jnxIpv4(12)	Class 1	jnxIpv4AddrTable(1) . Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxRmon(13)	Class 3	jnxRmonAlarmTable(1) . Objects are exposed only for the default logical system.
jnxLdp(14)	Class 2	jnxLdpTrapVars(1) . All instances within a logical system are exposed. Data will not be segregated down to the routing instance level.

Table 13: MIB Support for Routing Instances (Juniper Networks MIBs) (continued)

Object	Support Class	Description/Notes
jnxCos(15) jnxCosIfqStatsTable(1) jnxCosFcTable(2) jnxCosFcIdTable(3) jnxCosQstatTable(4)	Class 3	Objects are exposed only for the default logical system.
jnxScu(16) jnxScuStatsTable(1)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxRpf(17) jnxRpfStatsTable(1)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxCfgMgmt(18)	Class 3	Objects are exposed only for the default logical system.
jnxPMon(19) jnxPMonFlowTable(1) jnxPMonErrorTable(2) jnxPMonMemoryTable(3)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxSonet(20) jnxSonetAlarmTable(1)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxAtmCos(21) jnxCosAtmVcTable(1) jnxCosAtmScTable(2) jnxCosAtmVcQstatsTable(3) jnxCosAtmTrunkTable(4)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
ipSecFlowMonitorMIB(22)	—	—
jnxMac(23) jnxMacStats(1)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
apsMIB(24)	Class 3	Objects are exposed only for the default logical system.
jnxChassisDefines(25)	Class 3	Objects are exposed only for the default logical system.

Table 13: MIB Support for Routing Instances (Juniper Networks MIBs) (continued)

Object	Support Class	Description/Notes
jnxVpnMIB(26)	Class 2	All instances within a logical system are exposed. Data will not be segregated down to the routing instance level.
jnxSericesInfoMib(27)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxCollectorMIB(28)	Class 1	Only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.
jnxHistory(29)	—	—
jnxSpMIB(32)	Class 3	Objects are exposed only for the default logical system.

[Table 14 on page 160](#) shows Class 1 MIB objects (standard and enterprise-specific MIBs) supported by Junos OS. With Class 1 objects, only those logical interfaces (and their parent physical interfaces) that belong to a specific routing instance are exposed.

Table 14: Class 1 MIB Objects (Standard and Juniper MIBs)

Class	MIB	Objects
Class 1	802.3ad.mib	(dot3adAgg) MIB objects: dot3adAggTable dot3adAggPortListTable (dot3adAggPort) dot3adAggPortTable dot3adAggPortStatsTable dot3adAggPortDebugTable
	rfc2863a.mib	ifTable ifXTable ifStackTable
	rfc2011a.mib	ipAddrTable ipNetToMediaTable
	rtmib.mib	ipForward (ipCidrRouteTable)
	rfc2665a.mib	dot3StatsTable dot3ControlTable dot3PauseTable
	rfc2495a.mib	dsx1ConfigTable dsx1CurrentTable dsx1IntervalTable dsx1TotalTable dsx1FarEndCurrentTable dsx1FarEndIntervalTable dsx1FarEndTotalTable dsx1FracTable ...
	rfc2496a.mib	dsx3 (dsx3ConfigTable)
	rfc2115a.mib	frDlcmiTable (and related MIB objects)
	rfc3592.mib	sonetMediumTable (and related MIB objects)

Table 14: Class 1 MIB Objects (Standard and Juniper MIBs) (*continued*)

Class	MIB	Objects
	rfc3020.mib	mfrMIB mfrBundleTable mfrMibBundleLinkObjects mfrBundleIfIndexMappingTable (and related MIB objects)
	ospf2mib.mib	All objects
	ospf2trap.mib	All objects
	bgpmib.mib	All objects
	rfc2819a.mib	Example: etherStatsTable

Table 14: Class 1 MIB Objects (Standard and Juniper MIBs) (*continued*)

Class	MIB	Objects
Class 1	rfc2863a.mib	Examples: ifXtable ifStackTable
	rfc2665a.mib	etherMIB
	rfc2515a.mib	atmMIB objects Examples: atmInterfaceConfTable atmVplTable atmVclTable
	rfc2465.mib	ip-v6mib Examples: ipv6IfTable ipv6AddrPrefixTable ipv6NetToMediaTable ipv6RouteTable
	rfc2787a.mib	vrrp mib
	rfc2932.mib	ipMRouteMIB ipMRouteStdMIB
	mroutemib.mib	ipMRoute1MIBObjects
	isismib.mib	isisMIB
	pimmib.mib	pimMIB
	msdpmib.mib	msdpmib
	jnx-if-extensions.mib	Examples: ifJnxTable ifChassisTable
	jnx-dcu.mib	jnxDCUs
	jnx-atm.mib	

Table 14: Class 1 MIB Objects (Standard and Juniper MIBs) (*continued*)

Class	MIB	Objects
		Examples: <code>jnxAtmIfTable</code> <code>jnxAtmVcTable</code> <code>jnxAtmVpTable</code>
	<code>jnx-ipv4.mib</code>	<code>jnxipv4</code> Example: <code>jnxIpv4AddrTable</code>
	<code>jnx-cos.mib</code>	Examples: <code>jnxCosIfqStatsTable</code> <code>jnxCosQstatTable</code>
	<code>jnx-scu.mib</code>	Example: <code>jnxScuStatsTable</code>
	<code>jnx-rpf.mib</code>	Example: <code>jnxRpfStatsTable</code>
	<code>jnx-pmon.mib</code>	Example: <code>jnxPMonFlowTable</code>
	<code>jnx-sonet.mib</code>	Example: <code>jnxSonetAlarmTable</code>
	<code>jnx-atm-cos.mib</code>	Examples: <code>jnxCosAtmVcTable</code> <code>jnxCosAtmVcScTable</code> <code>jnxCosAtmVcQstatsTable</code> <code>jnxCosAtmTrunkTable</code>
	<code>jnx-mac.mib</code>	Example: <code>jnxMacStatsTable</code>
	<code>jnx-services.mib</code>	Example: <code>jnxSvcFlowTableAggStatsTable</code>
Class 1	<code>jnx-coll.mib</code>	<code>jnxCollectorMIB</code> Examples: <code>jnxCollPicIfTable</code> <code>jnxCollFileEntry</code>

Table 15 on page 164 shows Class 2 MIB objects (standard and enterprise-specific MIBs) supported by Junos OS. With Class 2 objects, all instances within a logical system are exposed. Data will not be segregated down to the routing instance level.

Table 15: Class 2 MIB Objects (Standard and Juniper MIBs)

Class	MIB	Objects
Class 2	rfc3813.mib	mplsLsrStdMIB Examples: mplsInterfaceTable mplsInSegmentTable mplsOutSegmentTable mplsLabelStackTable mplsXCTable (and related MIB objects)
	igmpmib.mib	igmpStdMIB NOTE: The igmpmib.mib is the draft version of the IGMP Standard MIB in the experimental tree. Junos OS does not support the original IGMP Standard MIB.
	l3vpn.mib	mplsVpnMIB
	jnx-mpls.mib	Example: mplsLspList
	jnx-ldp.mib	jnxLdp Example: jnxLdpStatsTable
	jnx-vpn.mib	jnxVpnMIB
	jnx-bgp.mib	jnxBgpMIB
	jnx-bgp-mib2.mib	jnxBgpM2Experiment

Table 16 on page 165 shows Class 3 MIB objects (standard and enterprise-specific MIBs) supported by Junos OS. With Class 3, objects are exposed only for the default logical system.

Table 16: Class 3 MIB Objects (Standard and Juniper MIBs)

Class	MIB	Objects
Class 3	rfc2819a.mib	rmonEvents alarmTable logTable eventTable agentxMIB
	rfc2925a.mib	pingmib
	rfc2925b.mib	tracerouteMIB
	jnxchassis.mib	jnxBoxAnatomy
	jnx-chassis-alarm.mib	jnxAlarms
	jnx-ping.mib	jnxPingMIB
	jnx-traceroute.mib	jnxTraceRouteMIB
	jnx-rmon.mib	jnxRmonAlarmTable
	jnx-cos.mib	Example: jnxCosFcTable
	jnx-cfgmgmt.mib	Example: jnxCfgMgmt
	jnx-sonetaps.mib	apsMIBObjects
	jnx-sp.mib	jnxSpMIB
	ggsn.mib	ejnmobileipABmib
	rfc1907.mib	snmpModules
	snmpModules	Examples: snmpMIB snmpFrameworkMIB

Table 17 on page 166 shows Class 4 MIB objects (standard and enterprise-specific MIBs) supported by Junos OS. With Class 4 objects, data is not segregated by routing instance. All instances are exposed.

Table 17: Class 4 MIB Objects (Standard and Juniper MIBs)

Class	MIB	Objects
Class 4	system	Example: sysORTable
	rfc2011a.mib	ip (ipDefaultTTL, ipInReceives) icmp
	rfc2012a.mib	tcp tcpConnTable ipv6TcpConnTable
	rfc2013a.mib	udp udpTable ipv6UdpTable
	rfc2790a.mib	hrSystem
	rfc2287a.mib	sysApplOBJ
	jnx-firewall.mib	jnxFirewalls
	jnx-ipv6.mib	jnxIpv6

Related Documentation

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [Support Classes for MIB Objects on page 166](#)
- [SNMP Traps Supported for Routing Instances on page 167](#)

Support Classes for MIB Objects

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

When a routing instance is specified, all routing-related MIB objects return data maintained by the routing instance in the request. For all other MIB objects, the data returned is segregated according to that routing instance. For example, only those interfaces assigned to that routing instance (for example, the logical interfaces [ifls] as well as their corresponding physical interfaces [ifds]) are exposed by the SNMP agent. Similarly, objects with an unambiguous attachment to an interface (for example, **addresses**) are segregated as well.

For those objects where the attachment is ambiguous (for example, objects in **sysApplMIB**), no segregation is done and all instances are visible in all cases.

Another category of objects is visible only when no logical system is specified (only within the default logical system) regardless of the routing instance within the default logical system. Objects in this category are Chassis MIB objects, objects in the SNMP group, RMON alarm, event and log groups, Ping MIB objects, configuration management objects, and V3 objects.

In summary, to support routing instances, MIB objects fall into one of the following categories:

- Class 1—Data is segregated according to the routing instance in the request. This is the most granular of the segregation classes.
- Class 2—Data is segregated according to the logical system specified in the request. The same data is returned for all routing instances that belong to a particular logical system. Typically, this applies to routing table objects where it is difficult to extract routing instance information or where routing instances do not apply.
- Class 3—Data is exposed only for the default logical system. The same set of data is returned for all routing instances that belong to the default logical system. If you specify another logical system (not the default), no data is returned. Typically this class applies to objects implemented in subagents that do not monitor logical system changes and register their objects using only the default context (for example, Chassis MIB objects).
- Class 4—Data is not segregated by routing instance. The same data is returned for all routing instances. Typically, this applies to objects implemented in subagents that monitor logical system changes and register or deregister all their objects for each logical system change. Objects whose values cannot be segregated by routing instance fall into this class.

See “SNMP MIBs Supported for Routing Instances” on page 156 for a list of the objects associated with each class.

Related Documentation

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [SNMP Traps Supported for Routing Instances on page 167](#)

SNMP Traps Supported for Routing Instances

Supported Platforms [M Series, MX Series, PTX Series, T Series](#)

You can restrict the trap receivers from receiving traps that are not related to the logical system networks to which they belong. To do this, include the **logical-system-trap-filter** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
logical-system-trap-filter;
```

If the **logical-system-trap-filter** statement is not included in the SNMP configuration, all traps are forwarded to the configured routing instance destinations. However, even when this statement is configured, the trap receiver associated with the default routing instance will receive all SNMP traps.

When configured under the trap-group object, all v1 and v2c traps that apply to routing instances (or interfaces belonging to a routing instance) have the routing instance name encoded in the community string. The encoding is identical to that used in request PDUs.

For traps configured under the v3 framework, the routing instance name is carried in the context field when the v3 message processing model has been configured. For other message processing models (v1 or v2c), the routing instance name is not carried in the trap message header (and not encoded in the community string).

**Related
Documentation**

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [Support Classes for MIB Objects on page 166](#)
- [SNMP MIBs Supported for Routing Instances on page 156](#)

Identifying a Routing Instance

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

With this feature, routing instances are identified by either the context field in v3 requests or encoded in the community string in v1 or v2c requests.

When encoded in a community string, the routing instance name appears first and is separated from the actual community string by the @ character.

To avoid conflicts with valid community strings that contain the @ character, the community is parsed only if typical community string processing fails. For example, if a routing instance named **RI** is configured, an SNMP request with **RI@public** is processed within the context of the **RI** routing instance. Access control (views, source address restrictions, access privileges, and so on) is applied according to the actual community string (the set of data after the @ character—in this case **public**). However, if the community string **RI@public** is configured, the protocol data unit (PDU) is processed according to that community and the embedded routing instance name is ignored.

Logical systems perform a subset of the actions of a physical router and have their own unique routing tables, interfaces, policies, and routing instances. When a routing instance is defined within a logical system, the logical system name must be encoded along with the routing instance using a slash (/) to separate the two. For example, if the routing instance **RI** is configured within the logical system **LS**, that routing instance must be encoded within a community string as **LS/RI@public**. When a routing instance is configured outside a logical system (within the default logical system), no logical system name (or / character) is needed.

Also, when a logical system is created, a default routing instance (named **default**) is always created within the logical system. This name should be used when querying data for that routing instance (for example, **LS/default@public**). For v3 requests, the name **logical system/routing instance** should be identified directly in the context field.



NOTE: To identify a virtual LAN (VLAN) spanning-tree instance (VSTP on MX Series 3D Universal Edge Routers), specify the routing instance name followed by a double colon (::) and the VLAN ID. For example, to identify VSTP instance for VLAN 10 in the global default routing instance, include `default::10@public` in the `context` (SNMPv3) or `community` (SNMPv1 or v2) string.

Related Documentation

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [Enabling SNMP Access over Routing Instances on page 169](#)
- [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)

Enabling SNMP Access over Routing Instances

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

To enable SNMP managers in routing instances other than the default routing instance to access SNMP information, include the **routing-instance-access** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
routing-instance-access;
```

If this statement is not included in the SNMP configuration, SNMP managers from routing instances other than the default routing instance cannot access SNMP information.

Related Documentation

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [Identifying a Routing Instance on page 168](#)
- [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)
- [Configuring Access Lists for SNMP Access over Routing Instances on page 172](#)

Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

You can specify the routing instance along with the client information when you add a client to an SNMP community. To specify the routing instance to which a client belongs, include the **routing-instance** statement followed by the routing instance name and client information in the SNMP configuration.

The following example shows the configuration statement to add routing instance **test-ri** to SNMP community **community1**.



NOTE: Routing instances specified at the `[edit snmp community community-name]` hierarchy level are added to the default logical system in the community.

```
[edit snmp]
community community1 {
  clients {
    10.209.152.33/32;
  }
  routing-instance test-ri {
    clients {
      10.19.19.1/32;
    }
  }
}
```

If the routing instance is defined within a logical system, include the **routing-instance** statement at the `[edit snmp community community-name logical-system logical-system-name]` hierarchy level, as in the following example:

```
[edit snmp]
community community1 {
  clients {
    10.209.152.33/32;
  }
  logical-system test-LS {
    routing-instance test-ri {
      clients {
        10.19.19.1/32;
      }
    }
  }
}
```

Related Documentation

- [Understanding SNMP Support for Routing Instances on page 155](#)
- [Identifying a Routing Instance on page 168](#)
- [Enabling SNMP Access over Routing Instances on page 169](#)
- [Configuring Access Lists for SNMP Access over Routing Instances on page 172](#)
- [Example: Configuring Interface Settings for a Routing Instance on page 170](#)

Example: Configuring Interface Settings for a Routing Instance

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

This example shows an **802.3ad ae0** interface configuration allocated to a routing instance named **INFrtid**:

```
[edit chassis]
aggregated-devices {
  ethernet {
```

```

        device-count 5;
    }
}
[edit interfaces ae0]
vlan-tagging;
aggregated-ether-options {
    minimum-links 2;
    link-speed 100m;
}
unit 0 {
    vlan-id 100;
    family inet {
        address 10.1.0.1/24;
    }
}
[edit interfaces fe-1/1/0]
fastether-options {
    802.3ad ae0;
}
[edit interfaces fe-1/1/1]
fastether-options {
    802.3ad ae0;
}
[edit routing-instances]
INFrtid {
    instance-type virtual-router;
    interface fe-1/1/0.0;
    interface fe-1/1/1.0;
    interface fe-1/1/5.0;
    interface ae0.0;
    protocols {
        ospf {
            area 0.0.0.0 {
                interface all;
            }
        }
    }
}
}

```

The following **snmpwalk** command shows how to retrieve SNMP-related information from **router1** and the 802.3ae bundle interface belonging to routing instance **INFrtid** with the SNMP community **public**:

```

router# snmpwalk -Os router1 INFrtid@public dot3adAggTable
dot3adAggMACAddress.59 = 0:90:69:92:93:f0
dot3adAggMACAddress.65 = 0:90:69:92:93:f0
dot3adAggActorSystemPriority.59 = 0
dot3adAggActorSystemPriority.65 = 0
dot3adAggActorSystemID.59 = 0:0:0:0:0:0
dot3adAggActorSystemID.65 = 0:0:0:0:0:0
dot3adAggAggregateOrIndividual.59 = true(1)
dot3adAggAggregateOrIndividual.65 = true(1)
dot3adAggActorAdminKey.59 = 0
dot3adAggActorAdminKey.65 = 0
dot3adAggActorOperKey.59 = 0
dot3adAggActorOperKey.65 = 0

```

```
dot3adAggPartnerSystemID.59 = 0:0:0:0:0:0
dot3adAggPartnerSystemID.65 = 0:0:0:0:0:0
dot3adAggPartnerSystemPriority.59 = 0
dot3adAggPartnerSystemPriority.65 = 0
dot3adAggPartnerOperKey.59 = 0
dot3adAggPartnerOperKey.65 = 0
dot3adAggCollectorMaxDelay.59 = 0
dot3adAggCollectorMaxDelay.65 = 0
```

- Related Documentation**
- [Understanding SNMP Support for Routing Instances on page 155](#)
 - [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)

Configuring Access Lists for SNMP Access over Routing Instances

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

You can create and maintain access lists to manage access to SNMP information. Access list configuration enables you to allow or deny SNMP access to clients of a specific routing instance.

The following example shows how to create an access list:

```
[edit snmp]
routing-instance-access {
  access-list {
    ri1 restrict;
    ls1/default;
    ls1/ri2;
    ls1*;
  }
}
```

The configuration given in the example:

- Restricts clients in **ri1** from accessing SNMP information.
- Allows clients in **ls1/default**, **ls1/ri2**, and all other routing instances with names starting with **ls1** to access SNMP information.

You can use the wildcard character (*) to represent a string in the routing instance name.



NOTE: You cannot restrict the SNMP manager of the default routing instance from accessing SNMP information.

- Related Documentation**
- [Understanding SNMP Support for Routing Instances on page 155](#)
 - [Enabling SNMP Access over Routing Instances on page 169](#)
 - [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)

CHAPTER 9

Configuring SNMP Remote Operations

- [SNMP Remote Operations Overview on page 173](#)
- [Using the Ping MIB for Remote Monitoring Devices Running Junos OS on page 176](#)
- [Starting a Ping Test on page 177](#)
- [Monitoring a Running Ping Test on page 178](#)
- [Gathering Ping Test Results on page 181](#)
- [Stopping a Ping Test on page 182](#)
- [Interpreting Ping Variables on page 182](#)
- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
- [Starting a Traceroute Test on page 184](#)
- [Monitoring a Running Traceroute Test on page 185](#)
- [Monitoring Traceroute Test Completion on page 189](#)
- [Gathering Traceroute Test Results on page 190](#)
- [Stopping a Traceroute Test on page 191](#)
- [Interpreting Traceroute Variables on page 192](#)

SNMP Remote Operations Overview

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

A SNMP remote operation is any process on the router that can be controlled remotely using SNMP. Junos OS currently provides support for two SNMP remote operations: the Ping MIB and Traceroute MIB, defined in RFC 2925. Using these MIBs, an SNMP client in the network management system (NMS) can:

- Start a series of operations on a router
- Receive notification when the operations are complete
- Gather the results of each operation

Junos OS also provides extended functionality to these MIBs in the Juniper Networks enterprise-specific extensions [jnxPingMIB](#) and [jnxTraceRouteMIB](#). For more information about [jnxPingMIB](#) and [jnxTraceRouteMIB](#), see [PING MIB](#) and [Traceroute MIB](#).

This topic covers the following sections:

- [SNMP Remote Operation Requirements on page 174](#)
- [Setting SNMP Views on page 174](#)
- [Setting Trap Notification for Remote Operations on page 175](#)
- [Using Variable-Length String Indexes on page 175](#)
- [Enabling Logging on page 176](#)

SNMP Remote Operation Requirements

To use SNMP remote operations, you should be experienced with SNMP conventions. You must also configure Junos OS to allow the use of the remote operation MIBs.

Setting SNMP Views

All remote operation MIBs supported by Junos OS require that the SNMP clients have read-write privileges. The default SNMP configuration of Junos OS does not provide clients with a community string with such privileges.

To set read-write privileges for an SNMP community string, include the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
community community-name {
  authorization authorization;
  view view-name;
}
view view-name {
  oid object-identifier (include | exclude);
}
```

Example: Setting SNMP Views

To create a community named **remote-community** that grants SNMP clients read-write access to the Ping MIB, **jnxPing** MIB, Traceroute MIB, and **jnxTraceRoute** MIB, include the following statements at the **[edit snmp]** hierarchy level:

```
snmp {
  view remote-view {
    oid 1.3.6.1.2.1.80 include; # pingMIB
    oid 1.3.6.1.4.1.2636.3.7 include; # jnxPingMIB
    oid 1.3.6.1.2.1.81 include; # traceRouteMIB
    oid 1.3.6.1.4.1.2636.3.8 include; # jnxTraceRouteMIB
  }
  community remote-community {
    view remote-view;
    authorization read-write;
  }
}
```

For more information about the **community** statement, see [“Configuring SNMP Communities” on page 93](#) and [community \(SNMP\)](#).

For more information about the **view** statement, see [“Configuring MIB Views” on page 111](#), [view \(Associating a MIB View with a Community\)](#), and [view \(Configuring a MIB View\)](#).

Setting Trap Notification for Remote Operations

In addition to configuring the remote operations MIB for trap notification, you must also configure Junos OS. You must specify a target host for remote operations traps.

To configure trap notification for SNMP remote operations, include the **categories** and **targets** statements at the **[edit snmp trap-group group-name]** hierarchy level:

```
[edit snmp trap-group group-name]
  categories {
    category;
  }
  targets {
    address;
  }
}
```

Example: Setting Trap Notification for Remote Operations

Specify 172.17.12.213 as a target host for all remote operation traps:

```
snmp {
  trap-group remote-traps {
    categories remote-operations;
    targets {
      172.17.12.213;
    }
  }
}
```

For more information about trap groups, see [“Configuring SNMP Trap Groups” on page 106](#).

Using Variable-Length String Indexes

All tabular objects in the remote operations MIBs supported by Junos OS are indexed by two variables of type **SnmpAdminString**. For more information about **SnmpAdminString**, see RFC 2571.

Junos OS does not handle **SnmpAdminString** any differently from the octet string variable type. However, the indexes are defined as variable length. When a variable length string is used as an index, the length of the string must be included as part of the object identifier (OID).

Example: Set Variable-Length String Indexes

To reference the **pingCtlTargetAddress** variable of a row in **pingCtlTable** where **pingCtlOwnerIndex** is **bob** and **pingCtlTestName** is **test**, use the following object identifier (OID):

```
pingMIB.pingObjects.pingCtlTable.pingCtlEntry.pingCtlTargetAddress."bob"."test"
1.3.6.1.2.1.80.1.2.1.4.3.98.111.98.4.116.101.115.116
```

For more information about the definition of the Ping MIB, see RFC 2925.

Enabling Logging

The SNMP error code returned in response to SNMP requests can only provide a generic description of the problem. The error descriptions logged by the remote operations process can often provide more detailed information about the problem and help you to solve the problem faster. This logging is not enabled by default. To enable logging, include the **flag general** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit]
snmp {
  traceoptions {
    flag general;
  }
}
```

For more information about traceoptions, see “[Tracing SNMP Activity on a Device Running Junos OS](#)” on page 201.

If the remote operations process receives an SNMP request that it cannot accommodate, the error is logged in the **/var/log/rmopd** file. To monitor this log file, issue the **monitor start rmopd** command in operational mode of the command-line interface (CLI).

Related Documentation

- [Using the Ping MIB for Remote Monitoring Devices Running Junos OS on page 176](#)
- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)

Using the Ping MIB for Remote Monitoring Devices Running Junos OS

Supported Platforms **M Series, MX Series, PTX Series, SRX Series, T Series, vSRX**

A ping test is used to determine whether packets sent from the local host reach the designated host and are returned. If the designated host can be reached, the ping test provides the approximate round-trip time for the packets. Ping test results are stored in **pingResultsTable** and **pingProbeHistoryTable**.

RFC 2925 is the authoritative description of the Ping MIB in detail and provides the ASN.1 MIB definition of the Ping MIB.

Related Documentation

- [SNMP Remote Operations Overview on page 173](#)
- [Starting a Ping Test on page 177](#)
- [Monitoring a Running Ping Test on page 178](#)
- [Gathering Ping Test Results on page 181](#)
- [Stopping a Ping Test on page 182](#)
- [Interpreting Ping Variables on page 182](#)

Starting a Ping Test

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

Before you start a ping test, configure a Ping MIB view. This allows SNMP **Set** requests on **pingMIB**. To start a ping test, create a row in **pingCtlTable** and set **pingCtlAdminStatus** to **enabled**. The minimum information that must be specified before setting **pingCtlAdminStatus** to **enabled** is:

- **pingCtlOwnerIndexSnmpAdminString**
- **pingCtlTestNameSnmpAdminString**
- **pingCtlTargetAddressInetAddress**
- **pingCtlTargetAddressTypeInetAddressType**
- **pingCtlRowStatusRowStatus**

For all other values, defaults are chosen unless otherwise specified. **pingCtlOwnerIndex** and **pingCtlTestName** are used as the index, so their values are specified as part of the object identifier (OID). To create a row, set **pingCtlRowStatus** to **createAndWait** or **createAndGo** on a row that does not already exist. A value of **active** for **pingCtlRowStatus** indicates that all necessary information has been supplied and the test can begin; **pingCtlAdminStatus** can be set to **enabled**. An SNMP **Set** request that sets **pingCtlRowStatus** to **active** will fail if the necessary information in the row is not specified or is inconsistent. For information about how to configure a view, see [“Setting SNMP Views” on page 174](#).

There are two ways to start a ping test:

- [Using Multiple Set Protocol Data Units \(PDUs\) on page 177](#)
- [Using a Single Set PDU on page 177](#)

Using Multiple Set Protocol Data Units (PDUs)

You can use multiple **Set** request PDUs (multiple PDUs, with one or more varbinds each) and set the following variables in this order to start the test:

- **pingCtlRowStatus** to **createAndWait**
- All appropriate test variables
- **pingCtlRowStatus** to **active**

Junos OS now verifies that all necessary information to run a test has been specified.

- **pingCtlAdminStatus** to **enabled**

Using a Single Set PDU

You can use a single **Set** request PDU (one PDU, with multiple varbinds) to set the following variables to start the test:

- **pingCtlRowStatus** to **createAndGo**
- All appropriate test variables
- **pingCtlAdminStatus** to **enabled**

Monitoring a Running Ping Test

When **pingCtlAdminStatus** is successfully set to **enabled**, the following is done before the acknowledgment of the SNMP **Set** request is sent back to the client:

- **pingResultsEntry** is created if it does not already exist.
- **pingResultsOperStatus** transitions to **enabled**.

For more information, see the following sections:

- [pingResultsTable](#) on page 178
- [pingProbeHistoryTable](#) on page 179
- [Generating Traps](#) on page 180

pingResultsTable

While the test is running, **pingResultsEntry** keeps track of the status of the test. The value of **pingResultsOperStatus** is **enabled** while the test is running and **disabled** when it has stopped.

The value of **pingCtlAdminStatus** remains **enabled** until you set it to **disabled**. Thus, to get the status of the test, you must examine **pingResultsOperStatus**.

The **pingCtlFrequency** variable can be used to schedule many tests for one **pingCtlEntry**. After a test ends normally (you did not stop the test) and the **pingCtlFrequency** number of seconds has elapsed, the test is started again just as if you had set **pingCtlAdminStatus** to **enabled**. If you intervene at any time between repeated tests (you set **pingCtlAdminStatus** to **disabled** or **pingCtlRowStatus** to **notInService**), the repeat feature is disabled until another test is started and ends normally. A value of 0 for **pingCtlFrequency** indicates this repeat feature is not active.

pingResultsIpTgtAddr and **pingResultsIpTgtAddrType** are set to the value of the resolved destination address when the value of **pingCtlTargetAddressType** is **dns**. When a test starts successfully and **pingResultsOperStatus** transitions to **enabled**:

- **pingResultsIpTgtAddr** is set to **null-string**.
- **pingResultsIpTgtAddrType** is set to **unknown**.

pingResultsIpTgtAddr and **pingResultsIpTgtAddrType** are not set until **pingCtlTargetAddress** can be resolved to a numeric address. To retrieve these values, poll **pingResultsIpTgtAddrType** for any value other than **unknown** after successfully setting **pingCtlAdminStatus** to **enabled**.

At the start of a test, **pingResultsSentProbes** is initialized to 1 and the first probe is sent. **pingResultsSentProbes** increases by 1 each time a probe is sent.

As the test runs, every **pingCtlTimeOut** seconds, the following occur:

- **pingProbeHistoryStatus** for the corresponding **pingProbeHistoryEntry** in **pingProbeHistoryTable** is set to **requestTimedOut**.
- A **pingProbeFailed** trap is generated, if necessary.
- An attempt is made to send the next probe.



NOTE: No more than one outstanding probe exists for each test.

For every probe, you can receive one of the following results:

- The target host acknowledges the probe with a response.
- The probe times out; there is no response from the target host acknowledging the probe.
- The probe could not be sent.

Each probe result is recorded in **pingProbeHistoryTable**. For more information about **pingProbeHistoryTable**, see [“pingProbeHistoryTable” on page 179](#).

When a response is received from the target host acknowledging the current probe:

- **pingResultsProbeResponses** increases by 1.
- The following variables are updated:
 - **pingResultsMinRtt**—Minimum round-trip time
 - **pingResultsMaxRtt**—Maximum round-trip time
 - **pingResultsAverageRtt**—Average round-trip time
 - **pingResultsRttSumOfSquares**—Sum of squares of round-trip times
 - **pingResultsLastGoodProbe**—Timestamp of the last response



NOTE: Only probes that result in a response from the target host contribute to the calculation of the round-trip time (RTT) variables.

When a response to the last probe is received or the last probe has timed out, the test is complete.

pingProbeHistoryTable

An entry in **pingProbeHistoryTable** (**pingProbeHistoryEntry**) represents a probe result and is indexed by three variables:

- The first two variables, **pingCtlOwnerIndex** and **pingCtlTestName**, are the same ones used for **pingCtlTable**, which identifies the test.

- The third variable, **pingProbeHistoryIndex**, is a counter to uniquely identify each probe result.

The maximum number of **pingProbeHistoryTable** entries created for a given test is limited by **pingCtlMaxRows**. If **pingCtlMaxRows** is set to 0, no **pingProbeHistoryTable** entries are created for that test.

Each time a probe result is determined, a **pingProbeHistoryEntry** is created and added to **pingProbeHistoryTable**. **pingProbeHistoryIndex** of the new **pingProbeHistoryEntry** is 1 greater than the last **pingProbeHistoryEntry** added to **pingProbeHistoryTable** for that test. **pingProbeHistoryIndex** is set to 1 if this is the first entry in the table. The same test can be run multiple times, so this index keeps growing.

If **pingProbeHistoryIndex** of the last **pingProbeHistoryEntry** added is 0xFFFFFFFF, the next **pingProbeHistoryEntry** added has **pingProbeHistoryIndex** set to 1.

The following are recorded for each probe result:

- **pingProbeHistoryResponse**—Time to live (TTL)
- **pingProbeHistoryStatus**—What happened and why
- **pingProbeHistoryLastRC**—Return code (RC) value of ICMP packet
- **pingProbeHistoryTime**—Timestamp when probe result was determined

When a probe cannot be sent, **pingProbeHistoryResponse** is set to 0. When a probe times out, **pingProbeHistoryResponse** is set to the difference between the time when the probe was discovered to be timed out and the time when the probe was sent.

Generating Traps

For any trap to be generated, the appropriate bit of **pingCtlTrapGeneration** must be set. You must also configure a trap group to receive remote operations. A trap is generated under the following conditions:

- A **pingProbeFailed** trap is generated every time **pingCtlTrapProbeFailureFilter** number of consecutive probes fail during the test.
- A **pingTestFailed** trap is generated when the test completes and at least **pingCtlTrapTestFailureFilter** number of probes fail.
- A **pingTestCompleted** trap is generated when the test completes and fewer than **pingCtlTrapTestFailureFilter** probes fail.



NOTE: A probe is considered a failure when **pingProbeHistoryStatus** of the probe result is anything besides **responseReceived**.

For information about how to configure a trap group to receive remote operations, see [“Configuring SNMP Trap Groups” on page 106](#) and [“Example: Setting Trap Notification for Remote Operations” on page 175](#).

Gathering Ping Test Results

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

You can either poll **pingResultsOperStatus** to find out when the test is complete or request that a trap be sent when the test is complete. For more information about **pingResultsOperStatus**, see “[pingResultsTable](#)” on page 178. For more information about Ping MIB traps, see “[Generating Traps](#)” on page 180.

The statistics calculated and then stored in **pingResultsTable** include:

- **pingResultsMinRtt**—Minimum round-trip time
- **pingResultsMaxRtt**—Maximum round-trip time
- **pingResultsAverageRtt**—Average round-trip time
- **pingResultsProbeResponses**—Number of responses received
- **pingResultsSentProbes**—Number of attempts to send probes
- **pingResultsRttSumOfSquares**—Sum of squares of round-trip times
- **pingResultsLastGoodProbe**—Timestamp of the last response

You can also consult **pingProbeHistoryTable** for more detailed information about each probe. The index used for **pingProbeHistoryTable** starts at 1, goes to 0xFFFFFFFF, and wraps to 1 again.

For example, if **pingCtlProbeCount** is 15 and **pingCtlMaxRows** is 5, then upon completion of the first run of this test, **pingProbeHistoryTable** contains probes like those in [Table 18 on page 181](#).

Table 18: Results in pingProbeHistoryTable: After the First Ping Test

pingProbeHistoryIndex	Probe Result
11	Result of 11th probe from run 1
12	Result of 12th probe from run 1
13	Result of 13th probe from run 1
14	Result of 14th probe from run 1
15	Result of 15th probe from run 1

Upon completion of the first probe of the second run of this test, **pingProbeHistoryTable** will contain probes like those in [Table 19 on page 182](#).

Table 19: Results in pingProbeHistoryTable: After the First Probe of the Second Test

pingProbeHistoryIndex	Probe Result
12	Result of 12th probe from run 1
13	Result of 13th probe from run 1
14	Result of 14th probe from run 1
15	Result of 15th probe from run 1
16	Result of 1st probe from run 2

Upon completion of the second run of this test, **pingProbeHistoryTable** will contain probes like those in [Table 20 on page 182](#).

Table 20: Results in pingProbeHistoryTable: After the Second Ping Test

pingProbeHistoryIndex	Probe Result
26	Result of 11th probe from run 2
27	Result of 12th probe from run 2
28	Result of 13th probe from run 2
29	Result of 14th probe from run 2
30	Result of 15th probe from run 2

History entries can be deleted from the MIB in two ways:

- More history entries for a given test are added and the number of history entries exceeds **pingCtlMaxRows**. The oldest history entries are deleted to make room for the new ones.
- You delete the entire test by setting **pingCtlRowStatus** to **destroy**.

Stopping a Ping Test

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

To stop an active test, set **pingCtlAdminStatus** to **disabled**. To stop the test and remove its **pingCtlEntry**, **pingResultsEntry**, and any **pingHistoryEntry** objects from the MIB, set **pingCtlRowStatus** to **destroy**.

Interpreting Ping Variables

Supported Platforms ACX Series, M Series, MX Series, PTX Series, T Series

This section clarifies the ranges for the following variables that are not explicitly specified in the Ping MIB:

- **pingCtlDataSize**—The value of this variable represents the total size of the payload (in bytes) of an outgoing probe packet. This payload includes the timestamp (8 bytes) that is used to time the probe. This is consistent with the definition of **pingCtlDataSize** (maximum value of 65,507) and the standard ping application.

If the value of **pingCtlDataSize** is between 0 and 8 inclusive, it is ignored and the payload is 8 bytes (the timestamp). The Ping MIB assumes all probes are timed, so the payload must always include the timestamp.

For example, if you wish to add an additional 4 bytes of payload to the packet, you must set **pingCtlDataSize** to 12.

- **pingCtlDataFill**—The first 8 bytes of the data segment of the packet is for the timestamp. After that, the **pingCtlDataFill** pattern is used in repetition. The default pattern (when **pingCtlDataFill** is not specified) is (00, 01, 02, 03 ... FF, 00, 01, 02, 03 ... FF, ...).
- **pingCtlMaxRows**—The maximum value is 255.
- **pingMaxConcurrentRequests**—The maximum value is 500.
- **pingCtlTrapProbeFailureFilter** and **pingCtlTrapTestFailureFilter**—A value of 0 for **pingCtlTrapProbeFailureFilter** or **pingCtlTrapTestFailureFilter** is not well defined by the Ping MIB. If **pingCtlTrapProbeFailureFilter** is 0, **pingProbeFailed** traps will not be generated for the test under any circumstances. If **pingCtlTrapTestFailureFilter** is 0, **pingTestFailed** traps will not be generated for the test under any circumstances.

Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

A traceroute test approximates the path packets take from the local host to the remote host.

RFC 2925 is the authoritative description of the Traceroute MIB in detail and provides the ASN.1 MIB definition of the Traceroute MIB.

- Related Documentation**
- [SNMP Remote Operations Overview on page 173](#)
 - [Starting a Traceroute Test on page 184](#)
 - [Monitoring a Running Traceroute Test on page 185](#)
 - [Monitoring Traceroute Test Completion on page 189](#)
 - [Gathering Traceroute Test Results on page 190](#)
 - [Stopping a Traceroute Test on page 191](#)
 - [Interpreting Traceroute Variables on page 192](#)

Starting a Traceroute Test

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Before you start a traceroute test, configure a Traceroute MIB view. This allows SNMP **Set** requests on **tracerouteMIB**. To start a test, create a row in **traceRouteCtlTable** and set **traceRouteCtlAdminStatus** to **enabled**. You must specify at least the following before setting **traceRouteCtlAdminStatus** to **enabled**:

- **traceRouteCtlOwnerIndexSnmpAdminString**
- **traceRouteCtlTestNameSnmpAdminString**
- **traceRouteCtlTargetAddressInetAddress**
- **traceRouteCtlRowStatusRowStatus**

For all other values, defaults are chosen unless otherwise specified.

traceRouteCtlOwnerIndex and **traceRouteCtlTestName** are used as the index, so their values are specified as part of the OID. To create a row, set **traceRouteCtlRowStatus** to **createAndWait** or **createAndGo** on a row that does not already exist. A value of **active** for **traceRouteCtlRowStatus** indicates that all necessary information has been specified and the test can begin; **traceRouteCtlAdminStatus** can be set to **enabled**. An SNMP **Set** request that sets **traceRouteCtlRowStatus** to **active** will fail if the necessary information in the row is not specified or is inconsistent. For information about how to configure a view, see [“Setting SNMP Views” on page 174](#).

There are two ways to start a traceroute test:

- [Using Multiple Set PDUs on page 184](#)
- [Using a Single Set PDU on page 184](#)

Using Multiple Set PDUs

You can use multiple **Set** request PDUs (multiple PDUs, with one or more varbinds each) and set the following variables in this order to start the test:

- **traceRouteCtlRowStatus** to **createAndWait**
- All appropriate test variables
- **traceRouteCtlRowStatus** to **active**

The Junos OS now verifies that all necessary information to run a test has been specified.

- **traceRouteCtlAdminStatus** to **enabled**

Using a Single Set PDU

You can use a single **Set** request PDU (one PDU, with multiple varbinds) to set the following variables to start the test:

- **traceRouteCtlRowStatus** to **createAndGo**

- All appropriate test variables
- **traceRouteCtlAdminStatus** to enabled

Related Documentation

- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
- [Monitoring a Running Traceroute Test on page 185](#)
- [SNMP Remote Operations Overview on page 173](#)
- [Monitoring Traceroute Test Completion on page 189](#)
- [Gathering Traceroute Test Results on page 190](#)
- [Stopping a Traceroute Test on page 191](#)
- [Interpreting Traceroute Variables on page 192](#)

Monitoring a Running Traceroute Test

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

When **traceRouteCtlAdminStatus** is successfully set to **enabled**, the following is done before the acknowledgment of the SNMP **Set** request is sent back to the client:

- **traceRouteResultsEntry** is created if it does not already exist.
- **traceRouteResultsOperStatus** transitions to enabled.

For more information, see the following sections:

- [traceRouteResultsTable on page 185](#)
- [traceRouteProbeResultsTable on page 186](#)
- [traceRouteHopsTable on page 187](#)
- [Generating Traps on page 188](#)

traceRouteResultsTable

While the test is running, this **traceRouteResultsTable** keeps track of the status of the test. The value of **traceRouteResultsOperStatus** is **enabled** while the test is running and **disabled** when it has stopped.

The value of **traceRouteCtlAdminStatus** remains **enabled** until you set it to **disabled**. Thus, to get the status of the test, you must examine **traceRouteResultsOperStatus**.

The **traceRouteCtlFrequency** variable can be used to schedule many tests for one **traceRouteCtlEntry**. After a test ends normally (you did not stop the test) and **traceRouteCtlFrequency** number of seconds has elapsed, the test is started again just as if you had set **traceRouteCtlAdminStatus** to **enabled**. If you intervene at any time between repeated tests (you set **traceRouteCtlAdminStatus** to **disabled** or **traceRouteCtlRowStatus** to **notInService**), the repeat feature is **disabled** until another test is started and ends normally. A value of 0 for **traceRouteCtlFrequency** indicates this repeat feature is not active.

traceRouteResultsIpTgtAddr and **traceRouteResultsIpTgtAddrType** are set to the value of the resolved destination address when the value of **traceRouteCtlTargetAddressType** is **dns**. When a test starts successfully and **traceRouteResultsOperStatus** transitions to **enabled**:

- **traceRouteResultsIpTgtAddr** is set to null-string.
- **traceRouteResultsIpTgtAddrType** is set to unknown.

traceRouteResultsIpTgtAddr and **traceRouteResultsIpTgtAddrType** are not set until **traceRouteCtlTargetAddress** can be resolved to a numeric address. To retrieve these values, poll **traceRouteResultsIpTgtAddrType** for any value other than **unknown** after successfully setting **traceRouteCtlAdminStatus** to **enabled**.

At the start of a test, **traceRouteResultsCurHopCount** is initialized to **traceRouteCtlInitialTtl**, and **traceRouteResultsCurProbeCount** is initialized to 1. Each time a probe result is determined, **traceRouteResultsCurProbeCount** increases by 1. While the test is running, the value of **traceRouteResultsCurProbeCount** reflects the current outstanding probe for which results have not yet been determined.

The **traceRouteCtlProbesPerHop** number of probes is sent for each time-to-live (TTL) value. When the result of the last probe for the current hop is determined, provided that the current hop is not the destination hop, **traceRouteResultsCurHopCount** increases by 1, and **traceRouteResultsCurProbeCount** resets to 1.

At the start of a test, if this is the first time this test has been run for this **traceRouteCtlEntry**, **traceRouteResultsTestAttempts** and **traceRouteResultsTestSuccesses** are initialized to 0.

At the end of each test execution, **traceRouteResultsOperStatus** transitions to **disabled**, and **traceRouteResultsTestAttempts** increases by 1. If the test was successful in determining the full path to the target, **traceRouteResultsTestSuccesses** increases by 1, and **traceRouteResultsLastGoodPath** is set to the current time.

traceRouteProbeResultsTable

Each entry in **traceRouteProbeHistoryTable** is indexed by five variables:

- The first two variables, **traceRouteCtlOwnerIndex** and **traceRouteCtlTestName**, are the same ones used for **traceRouteCtlTable** and to identify the test.
- The third variable, **traceRouteProbeHistoryIndex**, is a counter, starting from 1 and wrapping at FFFFFFFF. The maximum number of entries is limited by **traceRouteCtlMaxRows**.
- The fourth variable, **traceRouteProbeHistoryHopIndex**, indicates which hop this probe is for (the actual time-to-live or TTL value). Thus, the first **traceRouteCtlProbesPerHop** number of entries created when a test starts have a value of **traceRouteCtlInitialTtl** for **traceRouteProbeHistoryHopIndex**.
- The fifth variable, **traceRouteProbeHistoryProbeIndex**, is the probe for the current hop. It ranges from 1 to **traceRouteCtlProbesPerHop**.

While a test is running, as soon as a probe result is determined, the next probe is sent. A maximum of **traceRouteCtlTimeout** seconds elapses before a probe is marked with status **requestTimedOut** and the next probe is sent. There is never more than one outstanding probe per traceroute test. Any probe result coming back after a probe times out is ignored.

Each probe can:

- Result in a response from a host acknowledging the probe
- Time out with no response from a host acknowledging the probe
- Fail to be sent

Each probe status is recorded in **traceRouteProbeHistoryTable** with **traceRouteProbeHistoryStatus** set accordingly.

Probes that result in a response from a host record the following data:

- **traceRouteProbeHistoryResponse**—Round-trip time (RTT)
- **traceRouteProbeHistoryHAddrType**—The type of HAddr (next argument)
- **traceRouteProbeHistoryHAddr**—The address of the hop

All probes, regardless of whether a response for the probe is received, have the following recorded:

- **traceRouteProbeHistoryStatus**—What happened and why
- **traceRouteProbeHistoryLastRC**—Return code (RC) value of the ICMP packet
- **traceRouteProbeHistoryTime**—Timestamp when the probe result was determined

When a probe cannot be sent, **traceRouteProbeHistoryResponse** is set to 0. When a probe times out, **traceRouteProbeHistoryResponse** is set to the difference between the time when the probe was discovered to be timed out and the time when the probe was sent.

traceRouteHopsTable

Entries in **traceRouteHopsTable** are indexed by three variables:

- The first two, **traceRouteCtlOwnerIndex** and **traceRouteCtlTestName**, are the same ones used for **traceRouteCtlTable** and identify the test.
- The third variable, **traceRouteHopsHopIndex**, indicates the current hop, which starts at 1 (not **traceRouteCtlInitialTtl**).

When a test starts, all entries in **traceRouteHopsTable** with the given **traceRouteCtlOwnerIndex** and **traceRouteCtlTestName** are deleted. Entries in this table are only created if **traceRouteCtlCreateHopsEntries** is set to **true**.

A new **traceRouteHopsEntry** is created each time the first probe result for a given TTL is determined. The new entry is created whether or not the first probe reaches a host. The value of **traceRouteHopsHopIndex** is increased by 1 for this new entry.



NOTE: Any `traceRouteHopsEntry` can lack a value for `traceRouteHopsIpTgtAddress` if there are no responses to the probes with the given TTL.

Each time a probe reaches a host, the IP address of that host is available in the probe result. If the value of `traceRouteHopsIpTgtAddress` of the current `traceRouteHopsEntry` is not set, then the value of `traceRouteHopsIpTgtAddress` is set to this IP address. If the value of `traceRouteHopsIpTgtAddress` of the current `traceRouteHopsEntry` is the same as the IP address, then the value does not change. If the value of `traceRouteHopsIpTgtAddress` of the current `traceRouteHopsEntry` is different from this IP address, indicating a path change, a new `traceRouteHopsEntry` is created with:

- `traceRouteHopsHopIndex` variable increased by 1
- `traceRouteHopsIpTgtAddress` set to the IP address



NOTE: A new entry for a test is added to `traceRouteHopsTable` each time a new TTL value is used or the path changes. Thus, the number of entries for a test may exceed the number of different TTL values used.

When a probe result is determined, the value `traceRouteHopsSentProbes` of the current `traceRouteHopsEntry` increases by 1. When a probe result is determined, and the probe reaches a host:

- The value `traceRouteHopsProbeResponses` of the current `traceRouteHopsEntry` is increased by 1.
- The following variables are updated:
 - `traceRouteResultsMinRtt`—Minimum round-trip time
 - `traceRouteResultsMaxRtt`—Maximum round-trip time
 - `traceRouteResultsAverageRtt`—Average round-trip time
 - `traceRouteResultsRttSumOfSquares`—Sum of squares of round-trip times
 - `traceRouteResultsLastGoodProbe`—Timestamp of the last response



NOTE: Only probes that reach a host affect the round-trip time values.

Generating Traps

For any trap to be generated, the appropriate bit of `traceRouteCtlTrapGeneration` must be set. You must also configure a trap group to receive remote operations. Traps are generated under the following conditions:

- **traceRouteHopsIpTgtAddress** of the current probe is different from the last probe with the same TTL value (**traceRoutePathChange**).
- A path to the target could not be determined (**traceRouteTestFailed**).

A path to the target was determined (**traceRouteTestCompleted**).

For information about how to configure a trap group to receive remote operations, see “Configuring SNMP Trap Groups” on page 106 and “Example: Setting Trap Notification for Remote Operations” on page 175.

- See Also**
- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
 - [SNMP Remote Operations Overview on page 173](#)
 - [Starting a Traceroute Test on page 184](#)
 - [Monitoring Traceroute Test Completion on page 189](#)
 - [Gathering Traceroute Test Results on page 190](#)
 - [Stopping a Traceroute Test on page 191](#)
 - [Interpreting Traceroute Variables on page 192](#)

Monitoring Traceroute Test Completion

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

When a test is complete, **traceRouteResultsOperStatus** transitions from **enabled** to **disabled**. This transition occurs in the following situations:

- The test ends successfully. A probe result indicates that the destination has been reached. In this case, the current hop is the last hop. The rest of the probes for this hop are sent. When the last probe result for the current hop is determined, the test ends.
- **traceRouteCtlMaxTtl** threshold is exceeded. The destination is never reached. The test ends after the number of probes with TTL value equal to **traceRouteCtlMaxttl** have been sent.
- **traceRouteCtlMaxFailures** threshold is exceeded. The number of consecutive probes that end with status **requestTimedOut** exceeds **traceRouteCtlMaxFailures**.
- You end the test. You set **traceRouteCtlAdminStatus** to **disabled** or delete the row by setting **traceRouteCtlRowStatus** to **destroy**.
- You misconfigured the traceroute test. A value or variable you specified in **traceRouteCtlTable** is incorrect and will not allow a single probe to be sent. Because of the nature of the data, this error could not be determined until the test was started; that is, until after **traceRouteResultsOperStatus** transitioned to **enabled**. When this occurs, one entry is added to **traceRouteProbeHistoryTable** with **traceRouteProbeHistoryStatus** set to the appropriate error code.

If **traceRouteCtlTrapGeneration** is set properly, either the **traceRouteTestFailed** or **traceRouteTestCompleted** trap is generated.

- Related Documentation**
- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
 - [Monitoring a Running Traceroute Test on page 185](#)
 - [SNMP Remote Operations Overview on page 173](#)
 - [Starting a Traceroute Test on page 184](#)
 - [Gathering Traceroute Test Results on page 190](#)
 - [Stopping a Traceroute Test on page 191](#)
 - [Interpreting Traceroute Variables on page 192](#)

Gathering Traceroute Test Results

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

You can either poll **traceRouteResultsOperStatus** to find out when the test is complete or request that a trap be sent when the test is complete. For more information about **traceResultsOperStatus**, see “[traceRouteResultsTable](#)” on page 185. For more information about Traceroute MIB traps, see the Generating Traps section in “[Monitoring a Running Traceroute Test](#)” on page 185.

Statistics are calculated on a per-hop basis and then stored in **traceRouteHopsTable**. They include the following for each hop:

- **traceRouteHopsIpTgtAddressType**—Address type of host at this hop
- **traceRouteHopsIpTgtAddress**—Address of host at this hop
- **traceRouteHopsMinRtt**—Minimum round-trip time
- **traceRouteHopsMaxRtt**—Maximum round-trip time
- **traceRouteHopsAverageRtt**—Average round-trip time
- **traceRouteHopsRttSumOfSquares**—Sum of squares of round-trip times
- **traceRouteHopsSentProbes**—Number of attempts to send probes
- **traceRouteHopsProbeResponses**—Number of responses received
- **traceRouteHopsLastGoodProbe**—Timestamp of last response

You can also consult **traceRouteProbeHistoryTable** for more detailed information about each probe. The index used for **traceRouteProbeHistoryTable** starts at 1, goes to 0xFFFFFFFF, and wraps to 1 again.

For example, assume the following:

- **traceRouteCtlMaxRows** is 10.
- **traceRouteCtlProbesPerHop** is 5.

- There are eight hops to the target (the target being number eight).
- Each probe sent results in a response from a host (the number of probes sent is not limited by **traceRouteCtlMaxFailures**).

In this test, 40 probes are sent. At the end of the test, **traceRouteProbeHistoryTable** would have a history of probes like those in [Table 21 on page 191](#).

Table 21: traceRouteProbeHistoryTable

HistoryIndex	HistoryHopIndex	HistoryProbeIndex
31	7	1
32	7	2
33	7	3
34	7	4
35	7	5
36	8	1
37	8	2
38	8	3
39	8	4
40	8	5

Related Documentation

- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
- [Monitoring a Running Traceroute Test on page 185](#)
- [SNMP Remote Operations Overview on page 173](#)
- [Starting a Traceroute Test on page 184](#)
- [Monitoring Traceroute Test Completion on page 189](#)
- [Stopping a Traceroute Test on page 191](#)
- [Interpreting Traceroute Variables on page 192](#)

Stopping a Traceroute Test

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

To stop an active test, set **traceRouteCtlAdminStatus** to **disabled**. To stop a test and remove its **traceRouteCtlEntry**, **traceRouteResultsEntry**, **traceRouteProbeHistoryEntry**,

and **traceRouteProbeHistoryEntry** objects from the MIB, set **traceRouteCtlRowStatus** to **destroy**.

- Related Documentation**
- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
 - [Monitoring a Running Traceroute Test on page 185](#)
 - [SNMP Remote Operations Overview on page 173](#)
 - [Starting a Traceroute Test on page 184](#)
 - [Monitoring Traceroute Test Completion on page 189](#)
 - [Gathering Traceroute Test Results on page 190](#)
 - [Interpreting Traceroute Variables on page 192](#)

Interpreting Traceroute Variables

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

This topic contains information about the ranges for the following variables that are not explicitly specified in the Traceroute MIB:

- **traceRouteCtlMaxRows**—The maximum value for **traceRouteCtlMaxRows** is 2550. This represents the maximum TTL (255) multiplied by the maximum for **traceRouteCtlProbesPerHop** (10). Therefore, the **traceRouteProbeHistoryTable** accommodates one complete test at the maximum values for one **traceRouteCtlEntry**. Usually, the maximum values are not used and the **traceRouteProbeHistoryTable** is able to accommodate the complete history for many tests for the same **traceRouteCtlEntry**.
- **traceRouteMaxConcurrentRequests**—The maximum value is 50. If a test is running, it has one outstanding probe. **traceRouteMaxConcurrentRequests** represents the maximum number of traceroute tests that have **traceRouteResultsOperStatus** with a value of **enabled**. Any attempt to start a test with **traceRouteMaxConcurrentRequests** tests running will result in the creation of one probe with **traceRouteProbeHistoryStatus** set to **maxConcurrentLimitReached** and that test will end immediately.
- **traceRouteCtlTable**—The maximum number of entries allowed in this table is 100. Any attempt to create a 101st entry will result in a **BAD_VALUE** message for SNMPv1 and a **RESOURCE_UNAVAILABLE** message for SNMPv2.

- Related Documentation**
- [Using the Traceroute MIB for Remote Monitoring Devices Running Junos OS on page 183](#)
 - [Monitoring a Running Traceroute Test on page 185](#)
 - [SNMP Remote Operations Overview on page 173](#)
 - [Starting a Traceroute Test on page 184](#)
 - [Monitoring Traceroute Test Completion on page 189](#)
 - [Gathering Traceroute Test Results on page 190](#)

- [Stopping a Traceroute Test on page 191](#)

CHAPTER 10

Tracing SNMP Activity

- [Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS on page 195](#)
- [Tracing SNMP Activity on a Device Running Junos OS on page 201](#)
- [Example: Tracing SNMP Activity on page 205](#)

Monitoring SNMP Activity and Tracking Problems That Affect SNMP Performance on a Device Running Junos OS

Supported Platforms [ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series, vSRX](#)

The following sections contain information about monitoring the SNMP activity on devices running the Junos OS and identifying problems that might impact the SNMP performance on devices running Junos OS:

- [Checking for MIB Objects Registered with the snmpd on page 195](#)
- [Tracking SNMP Activity on page 197](#)
- [Monitoring SNMP Statistics on page 199](#)
- [Checking CPU Utilization on page 199](#)
- [Checking Kernel and Packet Forwarding Engine Response on page 200](#)

Checking for MIB Objects Registered with the snmpd

For the SNMP process to be able to access data related to a MIB object, the MIB object must be registered with the snmpd. When an SNMP subagent comes online, it tries to register the associated MIB objects with the snmpd. The snmpd maintains a mapping of the objects and the subagents with which the objects are associated. However, the registration attempt fails occasionally, and the objects remain unregistered with the snmpd until the next time the subagent restarts and successfully registers the objects.

When a network management system polls for data related to objects that are not registered with the snmpd, the snmpd returns either a **noSuchName** error (for SNMPv1 objects) or a **noSuchObject** error (for SNMPv2 objects).

You can use the following commands to check for MIB objects that are registered with the `snmpd`:

- **show snmp registered-objects**—Creates a `/var/log/snmp_reg_objs` file that contains the list of registered objects and their mapping to various subagents.
- **file show /var/log/snmp_reg_objs**—Displays the contents of the `/var/log/snmp_reg_objs` file.

The following example shows the steps for creating and displaying the `/var/log/snmp_reg_objs` file:

```
user@host> show snmp registered-objects
user@host> file show /var/log/snmp_reg_objs
-----
Registered MIB Objects
root_name =
-----
.1.2.840.10006.300.43.1.1.1.1.2 (dot3adAggMACAddress) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.3 (dot3adAggActorSystemPriority) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.4 (dot3adAggActorSystemID) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.5 (dot3adAggAggregateOrIndividual) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.6 (dot3adAggActorAdminKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.7 (dot3adAggActorOperKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.8 (dot3adAggPartnerSystemID) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.9 (dot3adAggPartnerSystemPriority) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.10 (dot3adAggPartnerOperKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.1.1.11 (dot3adAggCollectorMaxDelay) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.1.2.1.1 (dot3adAggPortListPorts) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.2 (dot3adAggPortActorSystemPriority) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.3 (dot3adAggPortActorSystemID) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.4 (dot3adAggPortActorAdminKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.5 (dot3adAggPortActorOperKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.6 (dot3adAggPortPartnerAdminSystemPriority) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.7 (dot3adAggPortPartnerOperSystemPriority) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.8 (dot3adAggPortPartnerAdminSystemID) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.9 (dot3adAggPortPartnerOperSystemID) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.10 (dot3adAggPortPartnerAdminKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.11 (dot3adAggPortPartnerOperKey) (/var/run/mib2d-11)
.1.2.840.10006.300.43.1.2.1.1.12 (dot3adAggPortSelectedAggID) (/var/run/mib2d-11)
---(more)---
```



NOTE: The `/var/log/snmp_reg_objs` file contains only those objects that are associated with the Junos OS processes that are up and running and registered with the `snmpd`, at the time of executing the `show snmp registered-objects` command. If a MIB object related to a Junos OS process that is up and running is not shown in the list of registered objects, you might want to restart the software process to retry object registration with the `snmpd`.

Copyright © 2017, Juniper Networks, Inc.

Monitoring SNMP Statistics

The **show snmp statistics extensive** operational mode command provides you with an option to review SNMP traffic, including traps, on a device. Output for the **show snmp statistics extensive** command shows real-time values and can be used to monitor values such as throttle drops, currently active, max active, not found, time out, max latency, current queued, total queued, and overflows. You can identify slowness in SNMP responses by monitoring the currently active count, because a constant increase in the currently active count is directly linked to slow or no response to SNMP requests.

Sample Output for the show snmp statistics extensive Command

```
user@host> show snmp statistics extensive
SNMP statistics:
Input:
  Packets: 226656, Bad versions: 0, Bad community names: 0,
  Bad community uses: 0, ASN parse errors: 0,
  Too bigs: 0, No such names: 0, Bad values: 0,
  Read onlys: 0, General errors: 0,
  Total request varbinds: 1967606, Total set varbinds: 0,
  Get requests: 18478, Get nexts: 75794, Set requests: 0,
  Get responses: 0, Traps: 0,
  Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
  Throttle drops: 27084, Duplicate request drops: 0
V3 Input:
  Unknown security models: 0, Invalid messages: 0
  Unknown pdu handlers: 0, Unavailable contexts: 0
  Unknown contexts: 0, Unsupported security levels: 0
  Not in time windows: 0, Unknown user names: 0
  Unknown engine ids: 0, Wrong digests: 0, Decryption errors: 0
Output:
  Packets: 226537, Too bigs: 0, No such names: 0,
  Bad values: 0, General errors: 0,
  Get requests: 0, Get nexts: 0, Set requests: 0,
  Get responses: 226155, Traps: 382
SA Control Blocks:
  Total: 222984, Currently Active: 501, Max Active: 501,
  Not found: 0, Timed Out: 0, Max Latency: 25
SA Registration:
  Registers: 0, Deregisters: 0, Removes: 0
Trap Queue Stats:
  Current queued: 0, Total queued: 0, Discards: 0, Overflows: 0
Trap Throttle Stats:
  Current throttled: 0, Throttles needed: 0
Snmp Set Stats:
  Commit pending failures: 0, Config lock failures: 0
  Rpc failures: 0, Journal write failures: 0
  Mgd connect failures: 0, General commit failures: 0
```

Checking CPU Utilization

High CPU usage of the software processes that are being queried, such as `snmpd` or `mib2d`, is another factor that can lead to slow response or no response. You can use the **show system processes extensive** operational mode command to check the CPU usage levels of the Junos OS processes.

Sample Output of show system processes extensive Command

```
user@host> show system processes extensive
last pid: 1415; load averages: 0.00, 0.00, 0.00 up 0+02:20:54 10:26:25
117 processes: 2 running, 98 sleeping, 17 waiting
```

```
Mem: 180M Active, 54M Inact, 39M Wired, 195M Cache, 69M Buf, 272M Free
Swap: 1536M Total, 1536M Free
```

PID	USERNAME	THR	PRI	NICE	SIZE	RES	STATE	TIME	WCPU	COMMAND
11	root	1	171	52	0K	12K	RUN	132:09	95.21%	idle
1184	root	1	97	0	35580K	9324K	select	4:16	1.61%	chassisd
177	root	1	-8	0	0K	12K	mdwait	0:51	0.00%	md7
119	root	1	-8	0	0K	12K	mdwait	0:20	0.00%	md4
13	root	1	-20	-139	0K	12K	WAIT	0:16	0.00%	swi7: clock sio
1373	root	1	96	0	15008K	12712K	select	0:09	0.00%	snmpd
1371	root	1	96	0	9520K	5032K	select	0:08	0.00%	jdiometerd
12	root	1	-40	-159	0K	12K	WAIT	0:07	0.00%	swi2: net
1375	root	2	96	0	15016K	5812K	select	0:06	0.00%	pfed
49	root	1	-8	0	0K	12K	mdwait	0:05	0.00%	md0
1345	root	1	96	0	10088K	4480K	select	0:05	0.00%	l2ald
1181	root	1	96	0	1608K	908K	select	0:05	0.00%	bslockd
23	root	1	-68	-187	0K	12K	WAIT	0:04	0.00%	irq10: fxp1
30	root	1	171	52	0K	12K	pgzero	0:04	0.00%	pagezero
1344	root	1	4	0	39704K	11444K	kqread	0:03	0.00%	rpd
1205	root	1	96	0	3152K	912K	select	0:03	0.00%	license-check
1372	root	1	96	0	28364K	6696K	select	0:03	0.00%	dcd
1374	root	1	96	0	11764K	7632K	select	0:02	0.00%	mib2d
1405	user	1	96	0	15892K	11132K	select	0:02	0.00%	cli
139	root	1	-8	0	0K	12K	mdwait	0:02	0.00%	md5
22	root	1	-80	-199	0K	12K	WAIT	0:02	0.00%	irq9: cbb1 fxp0
1185	root	1	96	0	4472K	2036K	select	0:02	0.00%	alarmd
4	root	1	-8	0	0K	12K	-	0:02	0.00%	g_down
3	root	1	-8	0	0K	12K	-	0:02	0.00%	g_up
43	root	1	-16	0	0K	12K	psleep	0:02	0.00%	vmkmemdaemon
1377	root	1	96	0	3776K	2256K	select	0:01	0.00%	irsd
48	root	1	-16	0	0K	12K	-	0:01	0.00%	schedcpu
99	root	1	-8	0	0K	12K	mdwait	0:01	0.00%	md3
953	root	1	96	0	4168K	2428K	select	0:01	0.00%	eventd
1364	root	1	96	0	4872K	2808K	select	0:01	0.00%	cfmd
15	root	1	-16	0	0K	12K	-	0:01	0.00%	yarrow
1350	root	1	96	0	31580K	7248K	select	0:01	0.00%	cosd
1378	root	1	96	0	19776K	6292K	select	0:01	0.00%	lpdfd

...

Checking Kernel and Packet Forwarding Engine Response

As mentioned in “[Understanding SNMP Implementation in Junos OS](#)” on page 15, some SNMP MIB data are maintained by the kernel or Packet Forwarding Engine. For such data to be available for the network management system, the kernel has to provide the required information to the SNMP subagent in mib2d. A slow response from the kernel can cause a delay in mib2d returning the data to the network management system. Junos OS adds an entry in the mib2d log file every time that an interface takes more than 10,000 microseconds to respond to a request for interface statistics. You can use the **show log log-filename | grep “kernel response time”** command to find out the response time taken by the kernel.

Checking the Kernel Response Time

```

user@host> show log mib2d | grep "kernel response time"
Aug 17 22:39:37 == kernel response time for
COS_IPVPN_DEFAULT_OUTPUT-t1-7/3/0:10:27.0-o: 9.126471 sec, range
(0.000007, 11.000806)

Aug 17 22:39:53 == kernel response time for
COS_IPVPN_DEFAULT_INPUT-t1-7/2/0:5:15.0-i: 5.387321 sec, range
(0.000007, 11.000806)

Aug 17 22:39:53 == kernel response time for ct1-6/1/0:9:15: 0.695406
sec, range (0.000007, 11.000806)

Aug 17 22:40:04 == kernel response time for t1-6/3/0:6:19: 1.878542
sec, range (0.000007, 11.000806)

Aug 17 22:40:22 == kernel response time for lsq-7/0/0: 2.556592 sec,
range (0.000007, 11.000806)

```

Related Documentation

- [Understanding SNMP Implementation in Junos OS on page 15](#)
- [Configuring SNMP on Devices Running Junos OS on page 85](#)
- [Optimizing the Network Management System Configuration for the Best Results on page 81](#)
- [Configuring Options on Managed Devices for Better SNMP Response Time on page 83](#)
- *Managing Traps and Informs*
- *Using the Enterprise-Specific Utility MIB to Enhance SNMP Coverage*

Tracing SNMP Activity on a Device Running Junos OS

Supported Platforms [ACX Series](#), [EX4600](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [T Series](#)

SNMP tracing operations track activity for SNMP agents and record the information in log files. The logged error descriptions provide detailed information to help you solve problems faster.

By default, Junos OS does not trace any SNMP activity. If you include the **traceoptions** statement at the **[edit snmp]** hierarchy level, the default tracing behavior is:

- Important activities are logged in files located in the **/var/log** directory. Each log is named after the SNMP agent that generates it. Currently, the following log files are created in the **/var/log** directory when the **traceoptions** statement is used:
 - chassisd
 - craftd
 - ilmid
 - mib2d
 - rmopd

- serviced
- snmpd
- When a trace file named *filename* reaches its maximum size, it is renamed *filename.0*, then *filename.1*, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. (For more information about how log files are created, see the [System Log Explorer](#).)
- Log files can be accessed only by the user who configured the tracing operation.

You cannot change the directory (*/var/log*) in which trace files are located. However, you can customize the other trace file settings by including the following statements at the **[edit snmp]** hierarchy level:

```
[edit snmp]
traceoptions {
  file <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  memory-trace;
  no-remote-trace;
  no-default-memory-trace;
}
```

These statements are described in the following sections:

- [Configuring the Number and Size of SNMP Log Files on page 202](#)
- [Configuring Access to the Log File on page 203](#)
- [Configuring a Regular Expression for Lines to Be Logged on page 203](#)
- [Configuring the Trace Operations on page 203](#)

Configuring the Number and Size of SNMP Log Files

By default, when the trace file reaches 128 kilobytes (KB) in size, it is renamed *filename.0*, then *filename.1*, and so on, until there are three trace files. Then the oldest trace file (*filename.2*) is overwritten.

You can configure the limits on the number and size of trace files by including the following statements at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file files number size size;
```

For example, set the maximum file size to 2 MB, and the maximum number of files to 20. When the file that receives the output of the tracing operation (*filename*) reaches 2 MB, *filename* is renamed *filename.0*, and a new file called *filename* is created. When the new *filename* reaches 2 MB, *filename.0* is renamed *filename.1* and *filename* is renamed *filename.0*. This process repeats until there are 20 trace files. Then the oldest file (*filename.19*) is overwritten by the newest file (*filename.0*).

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

Configuring Access to the Log File

By default, log files can be accessed only by the user who configured the tracing operation.

To specify that any user can read all log files, include the **file world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file world-readable;
```

To explicitly set the default behavior, include the **file no-world-readable** statement at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
file no-world-readable;
```

Configuring a Regular Expression for Lines to Be Logged

By default, the trace operation output includes all lines relevant to the logged activities.

You can refine the output by including the **match** statement at the **[edit snmp traceoptions file filename]** hierarchy level and specifying a regular expression (regex) to be matched:

```
[edit snmp traceoptions]
file filename match regular-expression;
```

Configuring the Trace Operations

By default, only important activities are logged. You can specify which trace operations are to be logged by including the following **flag** statement (with one or more tracing flags) at the **[edit snmp traceoptions]** hierarchy level:

```
[edit snmp traceoptions]
flag {
  all;
  configuration;
  database;
  events;
  general;
  interface-stats;
  nonvolatile-sets;
  pdu;
  policy;
  protocol-timeouts;
  routing-socket;
  server;
  subagent;
  timer;
  varbind-error;
}
```

Table 22 on page 204 describes the meaning of the SNMP tracing flags.

Table 22: SNMP Tracing Flags

Flag	Description	Default Setting
all	Log all operations.	Off
configuration	Log reading of the configuration at the [edit snmp] hierarchy level.	Off
database	Log events involving storage and retrieval in the events database.	Off
events	Log important events.	Off
general	Log general events.	Off
interface-stats	Log physical and logical interface statistics.	Off
nonvolatile-set	Log nonvolatile SNMP set request handling.	Off
pdu	Log SNMP request and response packets.	Off
policy	Log policy processing.	Off
protocol-timeouts	Log SNMP response timeouts.	Off
routing-socket	Log routing socket calls.	Off
server	Log communication with processes that are generating events.	Off
subagent	Log subagent restarts.	Off
timer	Log internal timer events.	Off
varbind-error	Log variable binding errors.	Off

To display the end of the log for an agent, issue the **show log agentd | last** operational mode command:

```
[edit]
user@host# run show log agentd | last
```

where **agent** is the name of an SNMP agent.

Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)
- [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)
- [Example: Tracing SNMP Activity on page 205](#)
- [Configuring SNMP](#)

Example: Tracing SNMP Activity

Supported Platforms [M Series, MX Series, PTX Series, T Series](#)

Trace information about SNMP packets:

```
[edit]
snmp {
  traceoptions {
    file size 10k files 5;
    flag pdu;
    flag protocol-timeouts;
    flag varbind-error;
  }
}
```

- Related Documentation**
- [Configuring SNMP on a Device Running Junos OS](#)
 - [Tracing SNMP Activity on a Device Running Junos OS on page 201](#)
 - [Configuration Statements at the \[edit snmp\] Hierarchy Level on page 78](#)

CHAPTER 11

SNMP FAQs

- [Junos OS SNMP FAQ Overview on page 207](#)
- [Junos OS SNMP FAQs on page 208](#)

Junos OS SNMP FAQ Overview

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [SRX Series](#), [T Series](#)

This document presents the most frequently asked questions about the features and technologies used to implement SNMP services on Juniper Networks devices using the Junos operating system.

SNMP enables users to monitor network devices from a central location. Many network management systems (NMS) are based on SNMP, and support for this protocol is a key feature of most network devices.

Juniper Networks provides many different platforms that support SNMP on the Junos OS. The Junos OS includes an onboard SNMP agent that provides remote management applications with access to detailed information about the devices on the network.

A typical SNMP implementation contains three components:

- Managed devices – Such as routers and switches.
- SNMP agent – Process that resides on a managed device and communicates with the NMS.
- NMS – A combination of hardware and software used to monitor and administer the network; network device that runs SNMP manager software. Also referred to as an SNMP manager.

The SNMP agent exchanges network management information with the SNMP manager (NMS). The agent responds to requests for information and actions from the manager. The SNMP manager collects information about network connectivity, activity, and events by polling managed devices.

SNMP implementation in the Junos OS uses a master SNMP agent (known as an SNMP process or `snmpd`) that resides on the managed device. Various subagents reside on different modules of the Junos OS as well (such as the Routing Engine), and these subagents are managed by the `snmpd`.

Related Documentation

- [Junos OS SNMP FAQs on page 208](#)

Junos OS SNMP FAQs

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [SRX Series](#), [T Series](#)

This Frequently Asked Questions technology overview covers these SNMP-related areas:

- [Junos OS SNMP Support FAQs on page 208](#)
- [Junos OS MIBs FAQs on page 209](#)
- [Junos OS SNMP Configuration FAQs on page 216](#)
- [SNMPv3 FAQs on page 220](#)
- [SNMP Interaction with Juniper Networks Devices FAQs on page 222](#)
- [SNMP Traps and Informs FAQs on page 224](#)
- [Junos OS Dual Routing Engine Configuration FAQs on page 230](#)
- [SNMP Support for Routing Instances FAQs on page 231](#)
- [SNMP Counters FAQs on page 232](#)

Junos OS SNMP Support FAQs

This section presents frequently asked questions and answers related to SNMP support on Junos OS.

Which SNMP versions does Junos OS support?

Junos OS supports SNMP version 1 (SNMPv1), version 2 (SNMPv2c), and version 3 (SNMPv3). By default, SNMP is disabled on a Juniper Networks device.

Which ports (sockets) does SNMP use?

The default port for SNMP queries is port 161. The default port for SNMP traps and informs is port 162. The port used for SNMP traps and informs is configurable, and you can configure your system to use ports other than the default port 162. However, the SNMP listening port will remain the same; this is established on the RFC.

Is SNMP support different among the Junos OS platforms?

No, SNMP support is not different among the Junos OS platforms. SNMP configuration, interaction, and behavior are the same on any Junos OS device. The only difference that might occur across platforms is MIB support.

See also [SNMP MIB Explorer](#) for a list of MIBs that are supported across the Junos OS platforms.

Does Junos OS support the user-based security model (USM)?

Yes, Junos OS supports USM as part of its support for SNMPv3. SNMPv3 contains more security measures than previous versions of SNMP, including providing a defined USM.

SNMPv3 USM provides message security through data integrity, data origin authentication, message replay protection, and protection against disclosure of the message payload.

Does Junos OS support the view-based access control model (VACM)?

Yes, Junos OS supports VACM as part of its support for SNMPv3. SNMPv3 contains more security measures than previous versions of SNMP, including providing a defined VACM. SNMPv3 VACM determines whether a specific type of access (read or write) to the management information is allowed.

Does Junos OS support SNMP informs?

Yes, Junos OS supports SNMP informs as part of its support for SNMPv3. SNMP informs are confirmed notifications sent from SNMP agents to SNMP managers when significant events occur on a network device. When an SNMP manager receives an inform, it sends a response to the sender to verify receipt of the inform.

Can I provision or configure a device using SNMP on Junos OS?

No, provisioning or configuring a device using SNMP is not allowed on Junos OS.

**Related
Documentation****Junos OS MIBs FAQs**

This section presents frequently asked questions and answers related to Junos OS MIBs.

What is a MIB?

A management information base (MIB) is a table of definitions for managed objects in a network device. MIBs are used by SNMP to maintain standard definitions of all of the components and their operating conditions within a network device. Each object in the MIB has an identifying code called an object identifier (OID).

MIBs are either standard or enterprise-specific. Standard MIBs are created by the Internet Engineering Task Force (IETF) and documented in various RFCs. Enterprise-specific MIBs are developed and supported by a specific equipment manufacturer.

For a list of supported standard MIBs, see “[Standard SNMP MIBs Supported by Junos OS](#)” on page 32.

For a list of Juniper Networks enterprise-specific MIBs, see “[Enterprise-Specific SNMP MIBs Supported by Junos OS](#)” on page 21.

Do MIB files reside on the Junos OS devices?

No, MIB files do not reside on the Junos OS devices. You must download the MIB files from the Juniper Networks Technical Publications page for the required Junos OS release: http://www.juniper.net/techpubs/en_US/release-independent/junos/mibs/mibs.html.

How do I compile and load the Junos OS MIBs onto an SNMP manager or NMS?

For your network management systems (NMSs) to identify and understand the MIB objects used by Junos OS, you must first load the MIB files to your NMS using a MIB compiler. A MIB compiler is a utility that parses the MIB information, such as the MIB object names, IDs, and data types for the NMS.

You can download the Junos OS MIB package from the Enterprise-Specific MIBs and Traps section at

http://www.juniper.net/techpubs/en_US/release-independent/junos/mibs/mibs.html or <http://www.juniper.net/techpubs/software/junos/index.html>.

The Junos OS MIB package has two folders: **StandardMibs**, containing standard MIBs supported on Juniper Networks devices, and **JuniperMibs**, containing Juniper Networks enterprise-specific MIBs. You *must* have the required standard MIBs downloaded and decompressed before downloading any enterprise-specific MIBs. There might be dependencies that require a particular standard MIB to be present on the compiler before loading a particular enterprise-specific MIB.

The Junos OS MIB package is available in **.zip** and **.tar** formats. Download the format appropriate for your requirements.

Use the following steps to load MIB files for devices running Junos OS:

1. Navigate to the appropriate Juniper Networks software download page and locate the **Enterprise MIBs** link under the **Enterprise-Specific MIBs and Traps** section.



NOTE: Although the link is titled Enterprise MIBs, both standard MIBs and enterprise-specific MIBs are available for download from this location.

2. Click the **TAR** or **ZIP** link to download the Junos OS MIB package.
3. Decompress the file (**.tar** or **.zip**) using an appropriate utility.



NOTE: Some commonly used MIB compilers are preloaded with standard MIBs. You can skip Step 4 and Step 5 and proceed to Step 6 if you already have the standard MIBs loaded on your system.

4. Load the standard MIB files from the **StandardMibs** folder.

Load the files in the following order:

- a. mib-SNMPv2-SMI.txt
- b. mib-SNMPv2-TC.txt
- c. mib-IANAifType-MIB.txt
- d. mib-IANA-RTPROTO-MIB.txt
- e. mib-ietf1907.txt

- f. mib-rfc2011a.txt
- g. mib-rfc2012a.txt
- h. mib-rfc2013a.txt
- i. mib-rfc2863a.txt

5. Load any remaining standard MIB files.



NOTE: You must follow the order specified in this procedure, and ensure that all standard MIBs are loaded before you load the enterprise-specific MIBs. There might be dependencies that require a particular standard MIB to be present on the compiler before loading a particular enterprise-specific MIB. Dependencies are listed in the **IMPORT** section of the MIB file.

6. After loading the standard MIBs, load the Juniper Networks enterprise-specific SMI MIB, **mib-jnx-smi.txt**, and the following optional SMI MIBs based on your requirements:

- mib-jnx-exp.txt—(Recommended) for Juniper Networks experimental MIB objects
- mib-jnx-js-smi.txt—(Optional) for Juniper Security MIB tree objects
- mib-jnx-ex-smi.txt—(Optional) for EX Series Ethernet Switches

7. Load any remaining desired enterprise-specific MIBs from the **JuniperMibs** folder.



TIP: While loading a MIB file, if the compiler returns an error message indicating that any of the objects are undefined, open the MIB file using a text editor and ensure that all the MIB files listed in the **IMPORT** section are loaded on the compiler. If any of the MIB files listed in the **IMPORT** section are not loaded on the compiler, load the missing file or files first, then try to load the MIB file that failed.

The system might return an error if files are not loaded in a particular order.

What is SMI?

Structure of Management Information Version (SMI) is a subset of Abstract Syntax Notation One (ASN.1), which describes the structure of objects. SMI is the notation syntax, or “grammar”, that is the standard for writing MIBs.

Which versions of SMI does Junos OS support?

The Junos OS supports SMIv1 for SNMPv1 MIBs, and SMIv2 for SNMPv2c and enterprise MIBs.

Does Junos OS support MIB II?

Yes, Junos OS supports MIB II, the second version of the MIB standard.

The features of MIB II include:

- Additions that reflect new operational requirements.
- Backward compatibility with the original MIBs and SNMP.
- Improved support for multiprotocol entities.
- Improved readability.

Refer to the relevant release documentation for a list of MIBs that are supported. Go to <http://www.juniper.net/techpubs/software/junos/index.html>.

Are the same MIBs supported across all Juniper Networks devices?

There are some common MIBs supported by all the Junos OS devices, such as the Interface MIB (ifTable), System MIB, and Chassis MIB. Some MIBs are supported only by functionalities on specific platforms. For example, the Bridge MIB is supported on the EX Series Ethernet Switches and the SRX Series Services Gateways for the branch.

What is the system object identifier (SYSOID) of a device? How do I determine the SYSOID of my device?

The jnx-chas-defines (Chassis Definitions for Router Model) MIB has a **jnxProductName** branch for every Junos OS device. The system object ID of a device is identical to the object ID of the **jnxProductName** for the platform. For example, for an M7i Multiservice Edge Router, the jnxProductNameM7i is .1.3.6.1.4.1.2636.1.1.1.2.10 in the jnxProductName branch, which is identical to the SYSOID of the M7i (.1.3.6.1.4.1.2636.1.1.1.2.10).

How can I determine if a MIB is supported on a platform? How can I determine which MIBs are supported by a device?

MIBs device and platform support is listed on the Junos OS Technical Documentation. See “[Enterprise-Specific SNMP MIBs Supported by Junos OS](#)” on page 21 and “[Standard SNMP MIBs Supported by Junos OS](#)” on page 32 documents to view the list of MIBs and supported Junos OS devices.

What can I do if the MIB OID query is not responding?

There can be various reasons why the MIB OID query stops responding. One reason could be that the MIB itself is unresponsive. To verify that the MIB responds, use the **show snmp mib walk | get MIB name | MIB OID** command:

- If the MIB responds, the communication issue exists between the SNMP master and SNMP agent. Possible reasons for this issue include network issues, an incorrect community configuration, an incorrect SNMP configuration, and so on.
- If the MIB does not respond, enable SNMP **traceoptions** to log PDUs and errors. All incoming and outgoing SNMP PDUs are logged. Check the **traceoptions** output to see if there are any errors.

If you continue to have problems with the MIB OID query, technical product support is available through the Juniper Networks Technical Assistance Center (JTAC).

What is the enterprise branch number for Junos OS?

The enterprise branch number for Junos OS is 2636. Enterprise branch numbers are used in SNMP MIB configurations, and they are also known as SMI network management private enterprise codes.

Which MIB displays the hardware and chassis details on a Juniper Networks device?

The Chassis MIB (`jnxchassis.mib`) displays the hardware and chassis details for each Juniper Networks device. It provides information about the router and its components. The Chassis MIB objects represent each component and its status.

Which MIB objects can I query to determine the CPU and memory utilization of the Routing Engine, Flexible PIC Concentrator (FPC), and PIC components on a device?

Query the Chassis MIB objects `jnxOperatingMemory`, `jnxOperatingtBuffer`, and `jnxOperatingCPU` to find out the CPU and memory utilization of the hardware components of a device.

Is the interface index (ifIndex) persistent?

The ifIndex is persistent when reboots occur if the Junos OS version remains the same, meaning the values assigned to the interfaces in the ifIndex do not change.

When there is a software upgrade, the device tries to keep the ifIndex persistent on a best effort basis. For Junos OS Release 10.0 and earlier, the ifIndex is not persistent when there is a software upgrade to Junos OS Release 10.1 and later.

Is it possible to set the ifAdminStatus?

SNMP is not allowed to set the ifAdminStatus.

Which MIB objects support SNMP set operations?

The Junos OS SNMP set operations are supported in the following MIB tables and variables:

- `snmpCommunityTable`
- `eventTable`
- `alarmTable`
- `snmpTargetAddrExtTable`
- `jnxPingCtlTable`
- `pingCtlTable`
- `traceRouteCtlTable`
- `jnxTraceRouteCtlTable`
- `sysContact.O`
- `sysName.O`
- `sysLocation.O`

- pingMaxConcurrentRequests.0
- traceRouteMaxConcurrentRequests.0
- usmUserSpinLock
- usmUserOwnAuthKeyChange
- usmUserPublic
- vacmSecurityToGroupTable (vacmGroupName, vacmSecurityToGroupStorageType, and vacmSecurityToGroupStatus)
- vacmAccessTable (vacmAccessContextMatch, vacmAccessReadViewName, vacmAccessWriteViewName, vacmAccessNotifyViewName, vacmAccessStorageType, and vacmAccessStatus)
- vacmViewSpinLock
- vacmViewTreeFamilyTable (vacmViewTreeFamilyMask, vacmViewTreeFamilyType, vacmViewTreeFamilyStorageType, and vacmViewTreeFamilyStatus)

Does Junos OS support remote monitoring (RMON)?

Yes, Junos OS supports RMON as defined in RFC 2819, *Remote Network Monitoring Management Information Base*. However, remote monitoring version 2 (RMON 2) is not supported.

Can I use SNMP to determine the health of the processes running on the Routing Engine?

Yes, you can use SNMP to determine the health of the Routing Engine processes by configuring the health monitoring feature. On Juniper Networks devices, RMON alarms and events provide much of the infrastructure needed to reduce the polling overhead from the NMS. However, you must set up the NMS to configure specific MIB objects into RMON alarms. This often requires device-specific expertise and customizing the monitoring application. Additionally, some MIB object instances that need monitoring are set only at initialization, or they change at runtime and cannot be configured in advance.

To address these issues, the health monitor extends the RMON alarm infrastructure to provide predefined monitoring for a selected set of object instances, such as file system usage, CPU usage, and memory usage, and includes support for unknown or dynamic object instances, such as Junos OS software processes.

To display the health monitoring configuration, use the **show snmp health-monitor** command:

```
user@host> show snmp health-monitor
interval 300;
rising-threshold 90;
falling-threshold 80;
```

When you configure the health monitor, monitoring information for certain object instances is available, as shown in [Table 23 on page 215](#).

Table 23: Monitored Object Instances

Object	Description
jnxHrStoragePercentUsed.1	Monitors the following file system on the router or switch: /dev/ad0s1a: This is the root file system mounted on /.
jnxHrStoragePercentUsed.2	Monitors the following file system on the router or switch: /dev/ad0s1e: This is the configuration file system mounted on /config.
jnxOperatingCPU (RE0)	Monitor CPU usage for Routing Engines RE0 and RE1. The index values assigned to the Routing Engines depend on whether the Chassis MIB uses a zero-based or a ones-based indexing scheme. Because the indexing scheme is configurable, the correct index is determined whenever the router is initialized and when there is a configuration change. If the router or switch has only one Routing Engine, the alarm entry monitoring RE1 is removed after five failed attempts to obtain the CPU value.
jnxOperatingCPU (RE1)	
jnxOperatingBuffer (RE0)	Monitor the amount of memory available on Routing Engines RE0 and RE1. Because the indexing of this object is identical to that used for jnxOperatingCPU, index values are adjusted depending on the indexing scheme used in the Chassis MIB. As with jnxOperatingCPU, the alarm entry monitoring RE1 is removed if the router or switch has only one Routing Engine.
jnxOperatingBuffer (RE1)	
sysApplElmtRunCPU	Monitors the CPU usage for each Junos OS software process. Multiple instances of the same process are monitored and indexed separately.
sysApplElmtRunMemory	Monitors the memory usage for each Junos OS software process. Multiple instances of the same process are monitored and indexed separately.

The system log entries generated for any health monitor events, such as thresholds crossed and errors, have a corresponding **HEALTHMONITOR** tag rather than a generic **SNMPD_RMON_EVENTLOG** tag. However, the health monitor sends generic **RMON risingThreshold** and **fallingThreshold** traps.

Are the Ping MIBs returned in decimal notation and ASCII?

Yes, both decimal notation and ASCII are supported, which is the standard implementation in SNMP. All strings are ASCII encoded.

The following example displays the Ping MIB in hexadecimal notation:

```
pingCtlTargetAddress.2.69.72.9.116.99.112.115.97.109.112.108.101 = 0a fa 01 02
```

This translates to ASCII:

```
pingCtlTargetAddress."EH"."tcpsample" = 0a fa 01 02
2= length of the string
69=E
72=H
9=length of second string
116=t
99 =c
112=p
```

```
115=s
97=a
109=m
112 =p
108 =l
101 =e
```

As of Junos OS Release 9.6 and later, the Junos OS CLI returns ASCII values using the command **show snmp mib get | get-next | walk ascii**.

The following example shows the output with the ASCII option:

```
user@host> show snmp mib walk pingCtlTargetAddress ascii
pingCtlTargetAddress."EH"."httpgetsampl" = http://www.yahoo.com
pingCtlTargetAddress."p1"."t2" = 74 c5 b3 06
pingCtlTargetAddress."p1"."t3" = 74 c5 b2 0c
```

The following example shows the output without the ASCII option:

```
user@host> show snmp mib walk pingCtlTargetAddress
pingCtlTargetAddress.2.69.72.13.104.116.116.112.103.101.116.115.97.109.112.108.101
= http://www.yahoo.com
pingCtlTargetAddress.2.112.49.2.116.50 = 74 c5 b3 06
pingCtlTargetAddress.2.112.49.2.116.51 = 74 c5 b2 0c
```

You can convert decimal and ASCII values using a decimal ASCII chart like the one at <http://www.asciichart.com>.

Is IPv6 supported by the Ping MIB for remote operations?

No, IPv6 is not supported.

Is there an SNMP MIB to show Address Resolution Protocol (ARP) table information? Are both IP and MAC addresses displayed in the same table?

Yes, the Junos OS supports the standard MIB **ipNetToMediaTable**, which is described in RFC 2011, *SNMPv2 Management Information Base for the Internet Protocol using SMIv2*. This table is used for mapping IP addresses to their corresponding MAC addresses.

Related Documentation

Junos OS SNMP Configuration FAQs

This section presents frequently asked questions and answers related to Junos OS SNMP configuration.

Can the Junos OS be configured for SNMPv1 and SNMPv3 simultaneously?

Yes, SNMP has backward compatibility, meaning that all three versions can be enabled simultaneously.

Can I filter specific SNMP queries on a device?

Yes, you can filter specific SNMP queries on a device using **exclude** and **include** statements.

The following example shows a configuration that blocks read-write operation on all OIDs under .1.3.6.1.2.1.1 for the community **test**:

```
user@host# show snmp
view system-exclude {
  oid .1.3.6.1.2.1.1 exclude;
  oid .1 include;
}
community test {
  view system-exclude;
  authorization read-write;
}
```

Can I change the SNMP agent engine ID?

Yes, the SNMP agent engine ID can be changed to the MAC address of the device, the IP address of the device, or any other desired value. Several examples are included here.

The following example shows how to use the MAC address of a device as the SNMP agent engine ID:

```
user@host# show snmp
engine-id {
  use-mac-address;
}
```

The following example shows how to use the IP address of a device as the SNMP agent engine ID:

```
user@host# show snmp
engine-id {
  use-default-ip-address;
}
```

The following example shows the use of a selected value, **AA** in this case, as the SNMP agent engine ID of a device:

```
user@host# show snmp
engine-id {
  local AA;
}
```

How can I configure a device with dual Routing Engines or a chassis cluster (SRX Series Services Gateways) for continued communication during a switchover?

When configuring for continued communication, the SNMP configuration should be identical between the Routing Engines. However, it is best to have separate Routing Engine IDs configured for each Routing Engine, especially when using SNMPv3.

The following example shows the configuration of the Routing Engines in a dual Routing Engine device. Notice that the Routing Engine IDs are set to the MAC addresses for each Routing Engine:

```
user@host# show groups
re0 {
  system {
```

```
    host-name PE3-re0;
  }
  interfaces {
    fxp0 {
      unit 0 {
        family inet {
          address 116.197.178.14/27;
          address 116.197.178.29/27 {
            master-only;
          }
        }
      }
    }
  }
}
snmp {
  engine-id {
    use-mac-address;
  }
}
}
rel {
  system {
    host-name PE3-rel;
  }
  interfaces {
    fxp0 {
      unit 0 {
        family inet {
          address 116.197.178.11/27;
          address 116.197.178.29/27 {
            master-only;
          }
        }
      }
    }
  }
}
snmp {
  engine-id {
    use-mac-address;
  }
}
}
```

The following is an example of an SNMPv3 configuration on a dual Routing Engine device:

```
user@host> show snmp name host1
v3 {
  vacm {
    security-to-group {
      security-model usm {
        security-name test123 {
          group test1;
        }
        security-name juniper {
          group test1;
        }
      }
    }
  }
}
```

```

    }
  }
  access {
    group test1 {
      default-context-prefix {
        security-model any {
          security-level authentication {
            read-view all;
          }
        }
      }
    }
    context-prefix MGMT_10 {
      security-model any {
        security-level authentication {
          read-view all;
        }
      }
    }
  }
}
target-address server1 {
  address 116.197.178.20;
  tag-list router1;
  routing-instance MGMT_10;
  target-parameters test;
}
target-parameters test {
  parameters {
    message-processing-model v3;
    security-model usm;
    security-level authentication;
    security-name juniper;
  }
  notify-filter filter1;
}
notify server {
  type trap;
  tag router1;
}
notify-filter filter1 {
  oid .1 include;
}
view all {
  oid .1 include;
}
community public {
  view all;
}
community comm1;
community comm2;
community comm3 {
  view all;
  authorization read-only;
  logical-system LDP-VPLS {
    routing-instance vpls-server1;
  }
}

```

```
    }  
  }  
  trap-group server1 {  
    targets {  
      116.197.179.22;  
    }  
  }  
  routing-instance-access;  
  traceoptions {  
    flag all;  
  }  
}
```

How can I track SNMP activities?

SNMP trace operations track activity of SNMP agents and record the information in log files.

A sample **traceoptions** configuration might look like this:

```
[edit snmp]  
user@host# set traceoptions flag all
```

When the **traceoptions flag all** statement is included at the **[edit snmp]** hierarchy level, the following log files are created:

- snmpd
- mib2d
- rmopd

Related Documentation

- [Junos OS SNMP Support FAQs on page 208](#)
- [Junos OS MIBs FAQs on page 209](#)
- [SNMPv3 FAQs on page 220](#)
- [SNMP Interaction with Juniper Networks Devices FAQs on page 222](#)
- [SNMP Traps and Informs FAQs on page 224](#)
- [SNMP Support for Routing Instances FAQs on page 231](#)
- [SNMP Counters FAQs on page 232](#)

SNMPv3 FAQs

This section presents frequently asked questions and answers related to SNMPv3.

Why is SNMPv3 important?

SNMP v3 provides enhanced security compared to the other versions of SNMP. It provides authentication and encryption of data. Enhanced security is important for managing devices at remote sites from the management stations.

In my system, the MIB object `snmpEngineBoots` is not in sync between two Routing Engines in a dual Routing Engine device. Is this normal behavior?

Yes, this is the expected behavior. Each Routing Engine runs its own SNMP process (`snmpd`), allowing each Routing Engine to maintain its own engine boots. However, if both routing engines have the same engine ID and the routing engine with lesser `snmpEngineBoots` value is selected as the master routing engine during the switchover process, the `snmpEngineBoots` value of the master routing engine is synchronized with the `snmpEngineBoots` value of the other routing engine.

Do I need the SNMP manager engine object identifier (OID) for informs?

Yes, the engine OID of the SNMP manager is required for authentication, and informs do not work without it.

I see the configuration of informs under the `[edit snmp v3]` hierarchy. Does this mean I cannot use informs with SNMPv2c?

Informs can be used with SNMPv2c. The following example shows the basic configuration for SNMPv3 informs on a device (note that the authentication and privacy is set to none):

```
[edit snmp]
v3 {
  usm {
    remote-engine 00000063000100a2c0a845b3 {
      user RU2_v3_sha_none {
        authentication-none;
        privacy-none;
      }
    }
  }
  vacm {
    security-to-group {
      security-model usm {
        security-name RU2_v3_sha_none {
          group g1_usm_auth;
        }
      }
    }
  }
  access {
    group g1_usm_auth {
      default-context-prefix {
        security-model usm {
          security-level authentication {
            read-view all;
            write-view all;
            notify-view all;
          }
        }
      }
    }
  }
  target-address TA2_v3_sha_none {
    address 192.168.69.179;
  }
}
```

```
tag-list tl1;
address-mask 255.255.252.0;
target-parameters TP2_v3_sha_none;
}
target-parameters TP2_v3_sha_none {
  parameters {
    message-processing-model v3;
    security-model usm;
    security-level none;
    security-name RU2_v3_sha_none;
  }
  notify-filter nf1;
}
notify N1_all_tl1_informs {
  type inform; # Replace "inform" with "trap" to convert informs to traps.
  tag tl1;
}
notify-filter nf1 {
  oid .1 include;
}
view all {
  oid .1 include;
}
}
```

You can convert the SNMPv3 informs to traps by setting the value of the **type** statement at the **[edit snmp v3 notify N1_all_tl1_informs]** hierarchy level to **trap** as shown in the following example:

```
user@host# set snmp v3 notify N1_all_tl1_informs type trap
```

Related Documentation

SNMP Interaction with Juniper Networks Devices FAQs

This section presents frequently asked questions and answers related to how SNMP interacts with Juniper Networks devices.

How frequently should a device be polled? What is a good polling rate?

It is difficult to give an absolute number for the rate of SNMP polls per second since the rate depends on the following two factors:

- The number of variable bindings in a protocol data unit (PDU)
- The response time for an interface from the Packet Forwarding Engine

In a normal scenario where no delay is being introduced by the Packet Forwarding Engine and there is one variable per PDU (a Get request), the response time is 130+ responses per second. However, with multiple variables in an SNMP request PDU (30 to 40 for GetBulk requests), the number of responses per second is much less. Because the Packet Forwarding Engine load can vary for each system, there is greater variation in how frequently a device should be polled.

Frequent polling of a large number of counters, especially statistics, can impact the device. We recommend the following optimization on the SNMP managers:

- Use the row-by-row polling method, not the column-by-column method.
- Reduce the number of variable bindings per PDU.
- Increase timeout values in polling and discovery intervals.
- Reduce the incoming packet rate at the SNMP process (snmpd).

For better SNMP response on the device, the Junos OS does the following:

- Filters out duplicate SNMP requests.
- Excludes interfaces that are slow in response from SNMP queries.

One way to determine a rate limit is to note an increase in the **Currently Active** count from the **show snmp statistics extensive** command.

The following is a sample output of the **show snmp statistics extensive** command:

```
user@host> show snmp statistics extensive
SNMP statistics:
  Input:
    Packets: 226656, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too big: 0, No such names: 0, Bad values: 0,
    Read only: 0, General errors: 0,
    Total request varbinds: 1967606, Total set varbinds: 0,
    Get requests: 18478, Get nexts: 75794, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 27084, Duplicate request drops: 0
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 0
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 0, Wrong digests: 0, Decryption errors: 0
  Output:
    Packets: 226537, Too big: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 226155, Traps: 382
  SA Control Blocks:
    Total: 222984, Currently Active: 501, Max Active: 501,
    Not found: 0, Timed Out: 0, Max Latency: 25
  SA Registration:
    Registers: 0, Deregisters: 0, Removes: 0
  Trap Queue Stats:
    Current queued: 0, Total queued: 0, Discards: 0, Overflows: 0
  Trap Throttle Stats:
    Current throttled: 0, Throttles needed: 0
  Snmp Set Stats:
    Commit pending failures: 0, Config lock failures: 0
    Rpc failures: 0, Journal write failures: 0
    Mgd connect failures: 0, General commit failures: 0
```

Does SNMP open dynamic UDP ports? Why?

The SNMP process opens two additional ports (sockets): one for IPv4 and one for IPv6. This enables the SNMP process to send traps.

I am unable to perform a MIB walk on the ifIndex. Why is this?

Any variable bindings or values with an access level of **not-accessible** cannot be queried directly because they are part of other variable bindings in the SNMP MIB table. The ifIndex has an access level of **not-accessible**. Therefore, it cannot be accessed directly because it is part of the variable bindings. However, the ifIndex can be accessed indirectly through the variable bindings.

I see SNMP_IPC_READ_ERROR messages when the SNMP process restarts on my system and also during Routing Engine switchover. Is this acceptable?

Yes, it is acceptable to see **SNMP_IPC_READ_ERROR** messages when the SNMP process is restarted, the system reboots, or during a Routing Engine switchover. If all the processes come up successfully and the SNMP operations are working properly, then these messages can be ignored.

What is the source IP address used in the response PDUs for SNMP requests? Can this be configured?

The source IP address used in the response PDUs for SNMP requests is the IP address of the outgoing interface to reach the destination. The source IP address cannot be configured for responses. It can only be configured for traps.

**Related
Documentation****SNMP Traps and Informs FAQs**

This section presents frequently asked questions and answers related to SNMP traps and informs.

Does the Junos OS impose any rate limiting on SNMP trap generation?

The Junos OS implements a trap-queuing mechanism to limit the number of traps that are generated and sent.

If a trap delivery fails, the trap is added back to the queue, and the delivery attempt counter and the next delivery attempt timer for the queue are reset. Subsequent attempts occur at progressive intervals of 1, 2, 4, and 8 minutes. The maximum delay between the attempts is 8 minutes, and the maximum number of attempts is 10. After 10 unsuccessful attempts, the destination queue and all traps in the queue are deleted.

Junos OS also has a throttle threshold mechanism to control the number of traps sent (default 500 traps) during a particular throttle interval (default 5 seconds). This helps ensure consistency in trap traffic, especially when a large number of traps are generated due to interface status changes.

The throttle interval begins when the first trap arrives at the throttle. All traps within the throttle threshold value are processed, and traps exceeding the threshold value are queued. The maximum size of all trap queues (the throttle queue and the destination queue) is 40,000 traps. The maximum size of any one queue is 20,000 traps. When a trap is added to the throttle queue, or if the throttle queue has exceeded the maximum size, the trap is moved to the top of the destination queue. Further attempts to send the trap from the destination queue are stopped for a 30-second period, after which the destination queue restarts sending the traps.



NOTE: For the Juniper Networks EX Series Ethernet Switch, the maximum size of all trap queues (the throttle queue and the destination queue) is 1,000 traps. The maximum size for any one queue on the EX Series is 500 traps.

I did not see a trap when I had a syslog entry with a critical severity. Is this normal? Can it be changed?

Not every syslog entry with critical severity is a trap. However, you can convert any syslog entry to a trap using the **event-options** statement.

The following example shows how to configure a **jnxSyslogTrap** whenever an **rpdp_ldap_nbrdown** syslog entry message error occurs.

```
user@host> show event-options
policy snmptrap {
  events rpd_ldap_nbrdown;
  then {
    raise-trap;
  }
}
```

Are SNMP traps compliant with the Alarm Reporting Function (X.733) on the Junos OS?

No, SNMP traps on the Junos OS are not X.733 compliant.

Can I set up filters for traps or informs?

Traps and informs can be filtered based on the trap category and the object identifier. You can specify categories of traps to receive per host by using the **categories** statement at the **[edit snmp trap-group trap-group]** hierarchy level. Use this option when you want to monitor only specific modules of the Junos OS.

The following example shows a sample configuration for receiving only **link**, **vrrp-events**, **services**, and **otn-alarms** traps:

```
[edit snmp]
trap-group jnpr {
  categories {
    link;
    vrrp-events;
    services;
    otn-alarms;
  }
}
```

```
targets {  
    192.168.69.179;  
}  
}
```

The Junos OS also has a more advanced filter option (**notify-filter**) for filtering specific traps or a group of traps based on their object identifiers.

The SNMPv3 configuration also supports filtering of SNMPv1 and SNMPv2 traps and excluding Juniper Networks enterprise-specific configuration management traps, as shown in the following configuration example:

```
[edit snmp]  
v3 {  
    vacm {  
        security-to-group {  
            security-model v2c {  
                security-name sn_v2c_trap {  
                    group gr_v2c_trap;  
                }  
            }  
        }  
    }  
    access {  
        group gr_v2c_trap {  
            default-context-prefix {  
                security-model v2c {  
                    security-level none {  
                        read-view all;  
                        notify-view all;  
                    }  
                }  
            }  
        }  
    }  
}  
target-address TA_v2c_trap {  
    address 10.209.196.166;  
    port 9001;  
    tag-list tg1;  
    target-parameters TP_v2c_trap;  
}  
target-parameters TP_v2c_trap {  
    parameters {  
        message-processing-model v2c;  
        security-model v2c;  
        security-level none;  
        security-name sn_v2c_trap;  
    }  
    notify-filter nf1;  
}  
notify v2c_notify {  
    type trap;  
    tag tg1;  
}  
notify-filter nf1 {  
    oid .1.3.6.1.4.1.2636.4.5 exclude;
```

```

oid .1 include;
}
snmp-community index1 {
  community-name "$9$tDLl01h7Nbw2axN"; ## SECRET-DATA
  security-name sn_v2c_trap;
  tag tg1;
}
view all {
  oid .1 include;
}
}

```

Can I simulate traps on a device?

Yes, you can use the **request snmp spoof-trap trap name** command for simulating a trap to the NMS that normally receives your device's traps. You can also add required values using the **variable-bindings** parameter.

The following example shows how to simulate a trap to the local NMS using variable bindings:

```

user@host> request snmp spoof-trap linkDown variable-bindings "ifIndex[116]=116,
ifAdminStatus[116]=1 ,ifOperStatus[116]=2 , ifName[116]=ge-1/0/1"

```

How do I generate a warm start SNMPv1 trap?

When the SNMP process is restarted under normal conditions, a warm start trap is generated if the system up time is more than 5 minutes. If the system up time is less than 5 minutes, a cold start trap is generated.

The NMS sees only the MIB OIDs and numbers, but not the names of the SNMP traps. Why?

Before the NMS can recognize the SNMP trap details, such as the names of the traps, it must first compile and understand the MIBs and then parse the MIB OIDs.

In the Junos OS, how can I determine to which category a trap belongs?

For a list of common traps and their categories, see [“Enterprise-Specific SNMP Traps Supported by Junos OS” on page 57](#).

Can I configure a trap to include the source IP address?

Yes, you can configure the **source-address**, **routing-instance**, or **logical-instance** name for the source IP address using the **trap-options** command:

```

user@host> show snmp trap-options
source-address 10.1.1.1;

```

Can I create a custom trap?

Yes, you can use the **jnxEventTrap** event script to create customized traps as needed.

In the following example, a Junos OS operations (op) script is triggered when a **UI_COMMIT_NOT_CONFIRMED** event is received. The Junos OS op script matches the complete message of the event and generates an SNMP trap.

Example: Junos OS Op Script

```
version 1.0;

ns junos = "http://xml.juniper.net/junos/*/junos";
ns xnm = "http://xml.juniper.net/xnm/1.1/xnm";
ns jcs = "http://xml.juniper.net/junos/commit-scripts/1.0";

param $event;
param $message;

match / {

    /*
     * trapm utility wants the following characters in the value to be escaped
     * '[', ']', ' ', '=', and ','
     */
    var $event-escaped = {
        call escape-string($text = $event, $vec = '[] =,');
    }

    var $message-escaped = {
        call escape-string($text = $message, $vec = '[] =,');
    }

    <op-script-results> {
    var $rpc = <request-snmp-spoof-trap> {
        <trap> "jnxEventTrap";
        <variable-bindings> "jnxEventTrapDescr[0]='Event-Trap' , "
        _ "jnxEventAvAttribute[1]='event' , "
        _ "jnxEventAvValue[1]='" _ $event-escaped _ "' , "
        _ "jnxEventAvAttribute[2]='message' , "
        _ "jnxEventAvValue[1]='" _ $message-escaped _ "'";
    }

    var $res = jcs:invoke($rpc);
    }

    template escape-string ($text, $vec) {

        if (jcs:empty($vec)) {
            expr $text;

        } else {
            var $index = 1;
            var $from = substring($vec, $index, 1);
            var $changed-value = {
                call replace-string($text, $from) {
                    with $to = {
                        expr "\\\";
                        expr $from;
                    }
                }
            }

            call escape-string($text = $changed-value, $vec = substring($vec, $index
+ 1));
        }
    }
}
```

```

template replace-string ($text, $from, $to) {

    if (contains($text, $from)) {
        var $before = substring-before($text, $from);
        var $after = substring-after($text, $from);
        var $prefix = $before _ $to;

        expr $before;
        expr $to;
        call replace-string($text = $after, $from, $to);

    } else {
        expr $text;
    }
}

```

After creating your customized trap, you must configure a policy on your device to tell the device what actions to take after it receives the trap.

Here is an example of a configured policy under the **[edit event-options]** hierarchy:

```

[edit event-options]
user@host> show
policy trap-on-event {
    events UI_COMMIT_NOT_CONFIRMED;
    attributes-match {
        UI_COMMIT_NOT_CONFIRMED.message matches complete;
    }
    then {
        event-script ev-syslog-trap.junos-op {
            arguments {
                event UI_COMMIT_NOT_CONFIRMED;
                message "${$.message}";
            }
        }
    }
}

```

Can I disable link up and link down traps on interfaces?

Yes, link up and link down traps can be disabled in the interface configuration. To disable the traps, use the **no-traps** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** and **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]** hierarchies for physical and logical interfaces.

```
(traps | no-traps);
```

I see the link up traps on logical interfaces, but I do not see the link down traps. Is this normal behavior?

For Ethernet and ATM types of interfaces, Junos OS does not send link down traps for a logical interface if the physical interface is down to prevent flooding alarms for the same root cause. However, when the physical interface and logical interfaces come back up,

traps are sent indicating link up. This is because the physical interface coming up does not necessarily mean the logical interfaces are also coming up.

For SONET types of interfaces with PPP encapsulation, Junos OS does send link down traps for a logical interface if the physical interface is down. When the physical interface and logical interfaces come back up, traps are sent for both the physical and logical interfaces indicating link up.

For SONET types of interfaces with HDLC encapsulation, Junos OS does not send link down traps for a logical interface if the physical interface is down. When the physical interface and logical interfaces come back up, traps are sent for both the physical and logical interfaces indicating link up.

For channelize interfaces with PPP encapsulation, Junos OS does send link down traps for a logical interface if the physical interface is down. When the physical interface and logical interfaces come back up, traps are sent for both the physical and logical interfaces indicating link up.

For channelize interfaces with HDLC encapsulation, Junos OS does not send link down traps for a logical interface if the physical interface is down. When the physical interface and logical interfaces come back up, traps are sent for both the physical and logical interfaces indicating link up.

Related Documentation

Junos OS Dual Routing Engine Configuration FAQs

This section presents frequently asked questions and answers related to the configuration of dual Routing Engines.

The SNMP configuration should be identical between the Routing Engines when configuring for continued communication. However, we recommend having separate Routing Engine IDs configured for each Routing Engine, when using SNMPv3.

In my system, the MIB object `snmpEngineBoots` is not in sync between two Routing Engines in a dual Routing Engine device. Is this normal behavior?

Yes. This is the normal behavior. Each Routing Engine runs its own SNMP process (`snmpd`) agent, allowing each Routing Engine to maintain its own engine boots.

Is there a way to identify that an address belongs to RE0, RE1, or the master Routing Engine management interface (`fxp0`) by looking at an SNMP walk?

No. When you do an SNMP walk on the device, it only displays the master Routing Engine management interface address.

What is the best way to tell if the current IP address belongs to `fxp0` or a Routing Engine, from a CLI session?

Routing Engines are mapped with the **fxp0** interface. This means that when you query RE0, the ifTable reports the **fxp0** interface address of RE0 only. Similarly, if you query RE1, the ifTable reports the **fxp0** interface address of RE1 only.

When there is a failover, the master hostname is changed since the hostname belongs to the Routing Engine. Is this correct?

Yes. You can configure the same hostname or different hostnames. Either would work.

If only the master IP address is configured (for example, 192.168.2.5), and the **sysDescr.0** object has the same string configured on both of the Routing Engines, then even after a switchover, the **sysDescr.0** object returns the same value. The following sample shows the results you get by using the **snmpget** command:

```
bng-junos-pool02: /c/svivek/PR_BRANCH/src> snmpget -c jnpr -v2c 192.168.2.5
sysDescr.0 system.sysDescr.0 = foo
```

SNMP Support for Routing Instances FAQs

This section presents frequently asked questions and answers related to how SNMP supports routing instances.

Can the SNMP manager access data for routing instances?

Yes, the Junos OS enables SNMP managers for all routing instances to request and manage SNMP data related to the corresponding routing instances and logical system networks.

Two different routing instance behaviors can occur, depending on where the clients originate:

- Clients from routing instances other than the default can access MIB objects and perform SNMP operations only on the logical system networks to which they belong.
- Clients from the default routing instance can access information related to all routing instances and logical system networks.

Routing instances are identified by either the context field in SNMPv3 requests or encoded in the community string in SNMPv1 or SNMPv2c requests.

When encoded in a community string, the routing instance name appears first and is separated from the actual community string by the @ character.

To avoid conflicts with valid community strings that contain the @ character, the community is parsed only if typical community string processing fails. For example, if a routing instance named **RI** is configured, an SNMP request with **RI@public** is processed within the context of the **RI** routing instance. Access control (including views, source address restrictions, and access privileges) is applied according to the actual community string (the set of data after the @ character—in this case **public**). However, if the community string **RI@public** is configured, the PDU is processed according to that community, and the embedded routing instance name is ignored.

Logical systems perform a subset of the actions of a physical router and have their own unique routing tables, interfaces, policies, and routing instances. When a routing instance

is defined within a logical system, the logical system name must be encoded along with the routing instance using a slash (/) to separate the two. For example, if the routing instance **RI** is configured within the logical system **LS**, that routing instance must be encoded within a community string as **LS/RI@public**. When a routing instance is configured outside a logical system (within the default logical system), no logical system name, or / character, is needed.

Additionally, when a logical system is created, a default routing instance named **default** is always created within the logical system. This name should be used when querying data for that routing instance, for example **LS/default@public**. For SNMPv3 requests, the name *logical system/routing instance* should be identified directly in the context field.

Can I access a list of all routing instances on a device?

Yes, you can access a list of all the routing instances on a device using the `vacmContextName` object in the `SNMP-VIEW-BASED-ACM` MIB. In SNMP, each routing instance becomes a VACM context; this is why the routing instances appear in the `vacmContextName` object.

Can I access a default routing instance from a client in another logical router or routing instance?

No, the SNMP agent can only access data of the logical router to which it is connected.

Related Documentation

SNMP Counters FAQs

This section presents frequently asked questions and answers related to SNMP counters.

Which MIB should I use for interface counters?

Interface management over SNMP is based on two tables: the **ifTable** and its extension the **ifXTable**. Both are described in RFC 1213, *Management Information Base for Network Management of TCP/IP-based internets: MIB-II* and RFC 2233, *The Interfaces Group MIB using SMIv2*.

Interfaces can have several layers, depending on the media, and each sublayer is represented by a separate row in the table. The relationship between the higher layer and lower layers is described in the **ifStackTable**.

The **ifTable** defines 32-bit counters for inbound and outbound octets (`ifInOctets/ifOutOctets`), packets (`ifInUcastPkts/ifOutUcastPkts`, `ifInNUcastPkts/ifOutNUcastPkts`), errors, and discards.

The **ifXTable** provides similar 64-bit counters, also called high capacity (HC) counters, for inbound and outbound octets (`ifHCInOctets/ifHCOctets`) and inbound packets (`ifHCInUcastPkts`).

When should 64-bit counters be used?

It is always good to use 64-bit counters because they contain statistics for both low and high capacity components.

Are the SNMP counters ifInOctets and ifOutOctets the same as the command reference show interfaces statistics in and out counters?

Yes, these are the same, but only if SNMP is enabled when the router boots up. If you power on a Juniper Networks device and then enable SNMP, the SNMP counters start from 0. SNMP counters do not automatically receive their statistics from the **show** command output. Similarly, using the **clear statistics** command does not clear the statistics that the SNMP counters collected, which can cause a discrepancy in the data that is seen by both processes.

Do the SNMP counters ifInOctets and ifOutOctets include the framing overhead for Point-to-Point Protocol (PPP) and High-Level Data Link Control (HDLC)?

Yes.

**Related
Documentation**

PART 3

Remote Monitoring (RMON) with SNMP

- [RMON Overview on page 237](#)
- [Configuring RMON Alarms and Events on page 241](#)
- [Monitoring RMON Alarms and Events on page 249](#)
- [Using RMON to Monitor Network Service Quality on page 257](#)

CHAPTER 12

RMON Overview

- [Understanding RMON Alarms on page 237](#)
- [Understanding RMON Events on page 238](#)

Understanding RMON Alarms

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, SRX Series, T Series](#)

An RMON alarm identifies:

- A specific MIB object that is monitored.
- The frequency of sampling.
- The method of sampling.
- The thresholds against which the monitored values are compared.

An RMON alarm can also identify a specific **eventTable** entry to be triggered when a threshold is crossed.

Configuration and operational values are defined in **alarmTable** in RFC 2819. Additional operational values are defined in Juniper Networks enterprise-specific extensions to **alarmTable** (**jnxRmonAlarmTable**).

This topic covers the following sections:

- [alarmTable on page 237](#)
- [jnxRmonAlarmTable on page 238](#)

alarmTable

alarmTable in the RMON MIB allows you to monitor and poll the following:

- **alarmIndex**—The index value for **alarmTable** that identifies a specific entry.
- **alarmInterval**—The interval, in seconds, over which data is sampled and compared with the rising and falling thresholds.
- **alarmVariable**—The MIB variable that is monitored by the alarm entry.
- **alarmSampleType**—The method of sampling the selected variable and calculating the value to be compared against the thresholds.

- **alarmValue**—The value of the variable during the last sampling period. This value is compared with the rising and falling thresholds.
- **alarmStartupAlarm**—The alarm sent when the entry is first activated.
- **alarmRisingThreshold**—The upper threshold for the sampled variable.
- **alarmFallingThreshold**—The lower threshold for the sampled variable.
- **alarmRisingEventIndex**—The **eventTable** entry used when a rising threshold is crossed.
- **alarmFallingEventIndex**—The **eventTable** entry used when a falling threshold is crossed.
- **alarmStatus**—Method for adding and removing entries from the table. It can also be used to change the state of an entry to allow modifications.



NOTE: If this object is not set to **valid**, the associated event alarm does not take any action.

jnxRmonAlarmTable

The **jnxRmonAlarmTable** is a Juniper Networks enterprise-specific extension to **alarmTable**. It provides additional operational information and includes the following objects:

- **jnxRmonAlarmGetFailCnt**—The number of times the internal **Get** request for the variable monitored by this entry has failed.
- **jnxRmonAlarmGetFailTime**—The value of **sysUpTime** when an internal **Get** request for the variable monitored by this entry last failed.
- **jnxRmonAlarmGetFailReason**—The reason an internal **Get** request for the variable monitored by this entry last failed.
- **jnxRmonAlarmGetOkTime**—The value of **sysUpTime** when an internal **Get** request for the variable monitored by this entry succeeded and the entry left the **getFailure** state.
- **jnxRmonAlarmState**—The current state of this RMON alarm entry.

To view the Juniper Networks enterprise-specific extensions to the RMON Events and Alarms and Event MIB, see
http://www.juniper.net/techpubs/en_US/junos16.1/topics/reference/mibs/mib-jnx-rmon.txt.

Related Documentation

- [Understanding RMON Events on page 238](#)
- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
- [Using alarmTable to Monitor MIB Objects on page 249](#)

Understanding RMON Events

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#)

An RMON event allows you to log the crossing of thresholds of other MIB objects. It is defined in **eventTable** for the RMON MIB.

This section covers the following topics:

- [eventTable on page 239](#)

eventTable

eventTable contains the following objects:

- **eventIndex**—An index that uniquely identifies an entry in **eventTable**. Each entry defines one event that is generated when the appropriate conditions occur.
- **eventDescription**—A comment describing the event entry.
- **eventType**—Type of notification that the probe makes about this event.
- **eventCommunity**—Trap group used if an SNMP trap is to be sent. If **eventCommunity** is not configured, a trap is sent to each trap group configured with the **rmon-alarm** category.
- **eventLastTimeSent**—Value of **sysUpTime** when this event entry last generated an event.
- **eventOwner**—Any text string specified by the creating management application or the command-line interface (CLI). Typically, it is used to identify a network manager (or application) and can be used for fine access control between participating management applications.
- **eventStatus**—Status of this event entry.



NOTE: If this object is not set to valid, no action is taken by the associated event entry. When this object is set to valid, all previous log entries associated with this entry (if any) are deleted.

Related Documentation

- [Understanding RMON Alarms on page 237](#)
- [Configuring an RMON Event Entry and Its Attributes on page 246](#)

Configuring RMON Alarms and Events

- [Understanding RMON Alarms and Events Configuration on page 241](#)
- [Minimum RMON Alarm and Event Entry Configuration on page 242](#)
- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
- [Configuring an RMON Event Entry and Its Attributes on page 246](#)
- [Example: Configuring an RMON Alarm and Event Entry on page 247](#)

Understanding RMON Alarms and Events Configuration

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#)

Junos OS supports monitoring routers from remote devices. These values are measured against thresholds and trigger events when the thresholds are crossed. You configure remote monitoring (RMON) alarm and event entries to monitor the value of a MIB object.

To configure RMON alarm and event entries, you include statements at the **[edit snmp]** hierarchy level of the configuration:

```
[edit snmp]
rmon {
  alarm index {
    description text-description;
    falling-event-index index;
    falling-threshold integer;
    falling-threshold-interval seconds;
    interval seconds;
    rising-event-index index;
    rising-threshold integer;
    request-type (get-next-request | get-request | walk-request);
    sample-type (absolute-value | delta-value);
    startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm);
    syslog-subtag syslog-subtag;
    variable oid-variable;
  }
  event index {
    community community-name;
    description description;
    type type;
  }
}
```

- Related Documentation**
- [Understanding RMON Alarms on page 237](#)
 - [Understanding RMON Events on page 238](#)
 - [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
 - [Configuring an RMON Event Entry and Its Attributes on page 246](#)

Minimum RMON Alarm and Event Entry Configuration

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

To enable RMON on the router, you must configure an alarm entry and an event entry. To do this, include the following statements at the **[edit snmp rmon]** hierarchy level:

```
[edit snmp rmon]
alarm index {
  rising-event-index index;
  rising-threshold integer;
  sample-type type;
  variable oid-variable;
}
event index;
```

- Related Documentation**
- [Understanding RMON Alarms and Events Configuration on page 241](#)
 - [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
 - [Configuring an RMON Event Entry and Its Attributes on page 246](#)

Configuring an RMON Alarm Entry and Its Attributes

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

An alarm entry monitors the value of a MIB variable. You can configure how often the value is sampled, the type of sampling to perform, and what event to trigger if a threshold is crossed.

This section discusses the following topics:

- [Configuring the Alarm Entry on page 243](#)
- [Configuring the Description on page 243](#)
- [Configuring the Falling Event Index or Rising Event Index on page 243](#)
- [Configuring the Falling Threshold or Rising Threshold on page 244](#)
- [Configuring the Interval on page 244](#)
- [Configuring the Falling Threshold Interval on page 244](#)
- [Configuring the Request Type on page 245](#)
- [Configuring the Sample Type on page 245](#)
- [Configuring the Startup Alarm on page 245](#)

- [Configuring the System Log Tag on page 246](#)
- [Configuring the Variable on page 246](#)

Configuring the Alarm Entry

An alarm entry monitors the value of a MIB variable. The **rising-event-index**, **rising-threshold**, **sample-type**, and **variable** statements are mandatory. All other statements are optional.

To configure the alarm entry, include the **alarm** statement and specify an index at the **[edit snmp rmon]** hierarchy level:

```
[edit snmp rmon]
alarm index {
  description description;
  falling-event-index index;
  falling-threshold integer;
  falling-threshold-interval seconds;
  interval seconds;
  rising-event-index index;
  rising-threshold integer;
  sample-type (absolute-value | delta-value);
  startup-alarm (falling-alarm | rising alarm | rising-or-falling-alarm);
  variable oid-variable;
}
```

index is an integer that identifies an alarm or event entry.

Configuring the Description

The description is a text string that identifies the alarm entry.

To configure the description, include the **description** statement and a description of the alarm entry at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
description description;
```

Configuring the Falling Event Index or Rising Event Index

The falling event index identifies the event entry that is triggered when a falling threshold is crossed. The rising event index identifies the event entry that is triggered when a rising threshold is crossed.

To configure the falling event index or rising event index, include the **falling-event-index** or **rising-event-index** statement and specify an index at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
falling-event-index index;
rising-event-index index;
```

index can be from 0 through 65,535. The default for both the falling and rising event index is 0.

Configuring the Falling Threshold or Rising Threshold

The falling threshold is the lower threshold for the monitored variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold, and the associated startup alarm is equal to **falling-alarm** or **rising-or-falling-alarm**. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the rising threshold. You must specify the falling threshold as an integer. Its default is 20 percent less than the rising threshold.

By default, the rising threshold is 0. The rising threshold is the upper threshold for the monitored variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold, and the associated **startup-alarm** is equal to **rising-alarm** or **rising-or-falling-alarm**. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the falling threshold. You must specify the rising threshold as an integer.

To configure the falling threshold or rising threshold, include the **falling-threshold** or **rising-threshold** statement at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
falling-threshold integer;
rising-threshold integer;
```

integer can be a value from -2,147,483,647 through 2,147,483,647.

Configuring the Interval

The interval represents the period of time, in seconds, over which the monitored variable is sampled and compared with the rising and falling thresholds.

To configure the interval, include the **interval** statement and specify the number of seconds at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
interval seconds;
```

seconds can be a value from 1 through 2,147,483,647. The default is 60 seconds.

Configuring the Falling Threshold Interval

The falling threshold interval represents the interval between samples when the rising threshold is crossed. Once the alarm crosses the falling threshold, the regular sampling interval is used.



NOTE: You cannot configure the falling threshold interval for alarms that have the request type set to walk-request.

To configure the falling threshold interval, include the **falling-threshold interval** statement at the **[edit snmp rmon alarm index]** hierarchy level and specify the number of seconds:

```
[edit snmp rmon alarm index]
  falling-threshold-interval seconds;
```

seconds can be a value from 1 through 2,147,483,647. The default is 60 seconds.

Configuring the Request Type

By default an RMON alarm can monitor only one object instance (as specified in the configuration). You can configure a **request-type** statement to extend the scope of the RMON alarm to include all object instances belonging to a MIB branch or to include the next object instance after the instance specified in the configuration.

To configure the request type, include the **request-type** statement at the **[edit snmp rmon alarm index]** hierarchy level and specify **get-next-request**, **get-request**, or **walk-request**:

```
[edit snmp rmon alarm index]
  request-type (get-next-request | get-request | walk-request);
```

walk extends the RMON alarm configuration to all object instances belonging to a MIB branch. **next** extends the RMON alarm configuration to include the next object instance after the instance specified in the configuration.

Configuring the Sample Type

The sample type identifies the method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is **absolute-value**, the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is **delta-value**, the value of the selected variable at the last sample is subtracted from the current value, and the difference is compared with the thresholds.

To configure the sample type, include the **sample-type** statement and specify the type of sample at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
  sample-type (absolute-value | delta-value);
```

- **absolute-value**—Actual value of the selected variable is compared against the thresholds.
- **delta-value**—Difference between samples of the selected variable is compared against the thresholds.

Configuring the Startup Alarm

The startup alarm identifies the type of alarm that can be sent when this entry is first activated. You can specify it as **falling-alarm**, **rising-alarm**, or **rising-or-falling-alarm**.

To configure the startup alarm, include the **startup-alarm** statement and specify the type of alarm at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
  startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm);
```

- **falling-alarm**—Generated if the first sample after the alarm entry becomes active is less than or equal to the falling threshold.
- **rising-alarm**—Generated if the first sample after the alarm entry becomes active is greater than or equal to the rising threshold.
- **rising-or-falling-alarm**—Generated if the first sample after the alarm entry becomes active satisfies either of the corresponding thresholds.

The default is **rising-or-falling-alarm**.

Configuring the System Log Tag

The **syslog-subtag** statement specifies the tag to be added to the system log message. You can specify a string of not more than 80 uppercase characters as the system log tag.

To configure the system log tag, include the **syslog-subtag** statement at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
syslog-subtag syslog-subtag;
```

Configuring the Variable

The variable identifies the MIB object that is being monitored.

To configure the variable, include the **variable** statement and specify the object identifier or object name at the **[edit snmp rmon alarm index]** hierarchy level:

```
[edit snmp rmon alarm index]
variable oid-variable;
```

oid-variable is a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1) or MIB object name (for example, ifInOctets.1).

- See Also**
- [Understanding RMON Alarms and Events Configuration on page 241](#)
 - [Understanding RMON Alarms on page 237](#)
 - [Understanding RMON Events on page 238](#)
 - [Configuring an RMON Event Entry and Its Attributes on page 246](#)
 - [Example: Configuring an RMON Alarm and Event Entry on page 247](#)

Configuring an RMON Event Entry and Its Attributes

Supported Platforms **M Series, MX Series, PTX Series, SRX Series, T Series**

An event entry generates a notification for an alarm entry when its rising or falling threshold is crossed. You can configure the type of notification that is generated. To configure the event entry, include the **event** statement at the **[edit snmp rmon]** hierarchy level. All statements except the **event** statement are optional.

```
[edit snmp rmon]
```



```

event index {
  community community-name;
  description description;
  type type;
}

```

index identifies an entry event.

community-name is the trap group that is used when generating a trap. If that trap group has the **rmon-alarm** trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group. If nothing is configured, all the trap groups are examined, and traps are sent using each group with the **rmon-alarm** category set.

description is a text string that identifies the entry.

The *type* variable of an event entry specifies where the event is to be logged. You can specify the type as one of the following:

- **log**—Adds the event entry to the **logTable**.
- **log-and-trap**—Sends an SNMP trap and creates a log entry.
- **none**—Sends no notification.
- **snmptrap**—Sends an SNMP trap.

The default for the event entry type is **log-and-trap**.

Related Documentation

- [Understanding RMON Alarms and Events Configuration on page 241](#)
- [Understanding RMON Alarms on page 237](#)
- [Understanding RMON Events on page 238](#)
- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
- [Example: Configuring an RMON Alarm and Event Entry on page 247](#)

Example: Configuring an RMON Alarm and Event Entry

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Configure an RMON alarm and event entry:

```

[edit snmp]
rmon {
  alarm 100 {
    description "input traffic on fxp0";
    falling-event-index 100;
    falling-threshold 10000;
    interval 60;
    rising-event-index 100;
    rising-threshold 100000;
    sample-type delta-value;
    startup-alarm rising-or-falling-alarm;
  }
}

```

```
        variable ifInOctets.1;  
    }  
    event 100 {  
        community bedrock;  
        description "emergency events";  
        type log-and-trap;  
    }  
}
```

**Related
Documentation**

- [Understanding RMON Alarms and Events Configuration on page 241](#)
- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
- [Configuring an RMON Event Entry and Its Attributes on page 246](#)

Monitoring RMON Alarms and Events

- [Using alarmTable to Monitor MIB Objects on page 249](#)
- [Using eventTable to Log Alarms on page 252](#)

Using alarmTable to Monitor MIB Objects

Supported Platforms [M Series, MX Series, T Series](#)

To use **alarmTable** to monitor a MIB object, perform the following tasks:

- [Creating an Alarm Entry on page 249](#)
- [Configuring the Alarm MIB Objects on page 249](#)
- [Activating a New Row in alarmTable on page 252](#)
- [Modifying an Active Row in alarmTable on page 252](#)
- [Deactivating a Row in alarmTable on page 252](#)

Creating an Alarm Entry

To create an alarm entry, first create a new row in **alarmTable** using the **alarmStatus** object. For example, create alarm #1 using the UCD command-line utilities:

```
snmpset -Os -v2c router community alarmStatus.1 i createRequest
```

Configuring the Alarm MIB Objects

Once you have created the new row in **alarmTable**, configure the following Alarm MIB objects:



NOTE: Other than **alarmStatus**, you cannot modify any of the objects in the entry if the associated **alarmStatus** object is set to valid.

- [alarmInterval on page 250](#)
- [alarmVariable on page 250](#)
- [alarmSampleType on page 250](#)
- [alarmValue on page 250](#)

- [alarmStartupAlarm](#) on page 250
- [alarmRisingThreshold](#) on page 251
- [alarmFallingThreshold](#) on page 251
- [alarmOwner](#) on page 251
- [alarmRisingEventIndex](#) on page 251
- [alarmFallingEventIndex](#) on page 252

alarmInterval

The interval, in seconds, over which data is sampled and compared with the rising and falling thresholds. For example, to set **alarmInterval** for alarm #1 to 30 seconds, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmInterval.1 i 30
```

alarmVariable

The object identifier of the variable to be sampled. During a **Set** request, if the supplied variable name is not available in the selected MIB view, a **badValue** error is returned. If at any time the variable name of an established **alarmEntry** is no longer available in the selected MIB view, the probe changes the status of **alarmVariable** to invalid. For example, to identify **ifInOctets.61** as the variable to be monitored, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmVariable.1 o .1.3.6.1.2.1.2.2.1.10.61
```

alarmSampleType

The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is **absoluteValue**, the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is **deltaValue**, the value of the selected variable at the last sample is subtracted from the current value, and the difference is compared with the thresholds. For example, to set **alarmSampleType** for alarm #1 to **deltaValue**, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmSampleType.1 i deltaValue
```

alarmValue

The value of the variable during the last sampling period. This value is compared with the rising and falling thresholds. If the sample type is **deltaValue**, this value equals the difference between the samples at the beginning and end of the period. If the sample type is **absoluteValue**, this value equals the sampled value at the end of the period.

alarmStartupAlarm

An alarm that is sent when this entry is first set to valid. If the first sample after this entry becomes valid is greater than or equal to **risingThreshold**, and **alarmStartupAlarm** is equal to **risingAlarm** or **risingOrFallingAlarm**, then a single rising alarm is generated. If the first

sample after this entry becomes valid is less than or equal to **fallingThreshold** and **alarmStartupAlarm** is equal to **fallingAlarm** or **risingOrFallingAlarm**, then a single falling alarm is generated. For example, to set **alarmStartupAlarm** for alarm #1 to **risingOrFallingAlarm**, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmStartupAlarm.1 i risingOrFallingAlarm
```

alarmRisingThreshold

A threshold for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold, and the associated **alarmStartupAlarm** is equal to **risingAlarm** or **risingOrFallingAlarm**. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches **alarmFallingThreshold**. For example, to set **alarmRisingThreshold** for alarm #1 to 100000, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmRisingThreshold.1 i 100000
```

alarmFallingThreshold

A threshold for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold, and the associated **alarmStartupAlarm** is equal to **fallingAlarm** or **risingOrFallingAlarm**. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches **alarmRisingThreshold**. For example, to set **alarmFallingThreshold** for alarm #1 to 10000, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmFallingThreshold.1 i 10000
```

alarmOwner

Any text string specified by the creating management application or the command-line interface (CLI). Typically, it is used to identify a network manager (or application) and can be used for fine access control between participating management applications.

alarmRisingEventIndex

The index of the **eventEntry** object that is used when a rising threshold is crossed. If there is no corresponding entry in **eventTable**, then no association exists. If this value is zero, no associated event is generated because zero is not a valid event index. For example, to set **alarmRisingEventIndex** for alarm #1 to 10, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmRisingEventIndex.1 i 10
```

[alarmFallingEventIndex](#)

The index of the **eventEntry** object that is used when a falling threshold is crossed. If there is no corresponding entry in **eventTable**, then no association exists. If this value is zero, no associated event is generated because zero is not a valid event index. For example, to set **alarmFallingEventIndex** for alarm #1 to 10, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community alarmFallingEventIndex.1 i 10
```

Activating a New Row in alarmTable

To activate a new row in **alarmTable**, set **alarmStatus** to **valid** using an SNMP **Set** request:

```
snmpset -Os -v2c router community alarmStatus.1 i valid
```

Modifying an Active Row in alarmTable

To modify an active row, first set **alarmStatus** to **underCreation** using an SNMP **Set** request:

```
snmpset -Os -v2c router community alarmStatus.1 i underCreation
```

Then change the row contents using an SNMP **Set** request:

```
snmpset -Os -v2c router community alarmFallingThreshold.1 i 1000
```

Finally, activate the row by setting **alarmStatus** to **valid** using an SNMP **Set** request:

```
snmpset -Os -v2c router community alarmStatus.1 i valid
```

Deactivating a Row in alarmTable

To deactivate a row in **alarmTable**, set **alarmStatus** to **invalid** using an SNMP **Set** request:

```
snmpset -Os -v2c router community alarmStatus.1 i invalid
```

Related Documentation

- [Understanding RMON Alarms on page 237](#)
- [Understanding RMON Events on page 238](#)
- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)

Using eventTable to Log Alarms

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

To use **eventTable** to log alarms, perform the following tasks:

- [Creating an Event Entry on page 253](#)
- [Configuring the MIB Objects on page 253](#)

- [Activating a New Row in eventTable on page 254](#)
- [Deactivating a Row in eventTable on page 255](#)

Creating an Event Entry

The RMON **eventTable** controls the generation of notifications from the router. Notifications can be logs (entries to **logTable** and **syslogs**) or SNMP traps. Each event entry can be configured to generate any combination of these notifications (or no notification). When an event specifies that an SNMP trap is to be generated, the trap group that is used when sending the trap is specified by the value of the associated **eventCommunity** object. Consequently, the community in the trap message will match the value specified by **eventCommunity**. If nothing is configured for **eventCommunity**, a trap is sent using each trap group that has the **rmon-alarm** category configured.

Configuring the MIB Objects

Once you have created the new row in **eventTable**, set the following objects:



NOTE: The **eventType** object is required. All other objects are optional.

- [eventType on page 253](#)
- [eventCommunity on page 253](#)
- [eventOwner on page 254](#)
- [eventDescription on page 254](#)

eventType

The type of notification that the router generates when the event is triggered.

This object can be set to the following values:

- **log**—Adds the event entry to **logTable**.
- **log-and-trap**—Sends an SNMP trap and creates a log entry.
- **none**—Sends no notification.
- **snmptrap**—Sends an SNMP trap.

For example, to set **eventType** for event #1 to **log-and-trap**, use the following SNMP Set request:

```
snmpset -Os -v2c router community eventType.1 i log-and-trap
```

eventCommunity

The trap group that is used when generating a trap (if **eventType** is configured to send traps). If that trap group has the **rmon-alarm** trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of **eventCommunity**). If nothing is

configured, traps are sent to each group with the **rmon-alarm** category set. For example, to set **eventCommunity** for event #1 to **boy-elroy**, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community eventCommunity.1 s "boy-elroy"
```



NOTE: The **eventCommunity** object is optional. If you do not set this object, then the field is left blank.

eventOwner

Any text string specified by the creating management application or the command-line interface (CLI). Typically, it is used to identify a network manager (or application) and can be used for fine access control between participating management applications.

For example, to set **eventOwner** for event #1 to **george jetson**, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community eventOwner.1 s "george jetson"
```



NOTE: The **eventOwner** object is optional. If you do not set this object, then the field is left blank.

eventDescription

Any text string specified by the creating management application or the command-line interface (CLI). The use of this string is application dependent.

For example, to set **eventDescription** for event #1 to **spacelys sprockets**, use the following SNMP **Set** request:

```
snmpset -Os -v2c router community eventDescription.1 s "spacelys sprockets"
```



NOTE: The **eventDescription** object is optional. If you do not set this object, then the field is left blank.

Activating a New Row in eventTable

To activate the new row in **eventTable**, set **eventStatus** to **valid** using an SNMP **Set** request such as:

```
snmpset -Os -v2c router community eventStatus.1 i valid
```


Deactivating a Row in eventTable

To deactivate a row in **eventTable**, set **eventStatus** to **invalid** using an SNMP **Set** request such as:

```
snmpset -Os -v2c router community eventStatus.1 i invalid
```

Related Documentation

- [Understanding RMON Alarms on page 237](#)
- [Understanding RMON Events on page 238](#)
- [Configuring an RMON Event Entry and Its Attributes on page 246](#)

CHAPTER 15

Using RMON to Monitor Network Service Quality

- [Understanding RMON for Monitoring Service Quality on page 257](#)
- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 261](#)
- [Defining and Measuring Network Availability on page 262](#)
- [Measuring Health on page 268](#)
- [Measuring Performance on page 274](#)

Understanding RMON for Monitoring Service Quality

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, SRX Series, T Series](#)

Health and performance monitoring can benefit from the remote monitoring of SNMP variables by the local SNMP agents running on each router. The SNMP agents compare MIB values against predefined thresholds and generate exception alarms without the need for polling by a central SNMP management platform. This is an effective mechanism for proactive management, as long as the thresholds have baselines determined and set correctly. For more information, see RFC 2819, *Remote Network Monitoring MIB*.


This topic includes the following sections:

- [Setting Thresholds on page 257](#)
- [RMON Command-Line Interface on page 258](#)
- [RMON Event Table on page 259](#)
- [RMON Alarm Table on page 259](#)
- [Troubleshooting RMON on page 260](#)

Setting Thresholds

By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside of the allowable operational range. (See [Figure 3 on page 258](#).)

Figure 3: Setting Thresholds



Events are only generated when the threshold is first crossed in any one direction rather than after each sample period. For example, if a rising threshold crossing event is raised, no more threshold crossing events will occur until a corresponding falling event. This considerably reduces the quantity of alarms that are produced by the system, making it easier for operations staff to react when alarms do occur.

To configure remote monitoring, specify the following pieces of information:

- The variable to be monitored (by its SNMP object identifier)
- The length of time between each inspection
- A rising threshold
- A falling threshold
- A rising event
- A falling event

Before you can successfully configure remote monitoring, you should identify what variables need to be monitored and their allowable operational range. This requires some period of baselining to determine the allowable operational ranges. An initial baseline period of at least three months is not unusual when first identifying the operational ranges and defining thresholds, but baseline monitoring should continue over the life span of each monitored variable.

RMON Command-Line Interface

Junos OS provides two mechanisms you use to control the Remote Monitoring agent on the router: command-line interface (CLI) and SNMP. To configure an RMON entry using the CLI, include the following statements at the **[edit snmp]** hierarchy level:

```
rmon {
  alarm index {
    description;
    falling-event-index;
    falling-threshold;
    intervals;
    rising-event-index;
```

```

    rising-threshold;
    sample-type (absolute-value | delta-value);
    startup-alarm (falling | rising | rising-or-falling);
    variable;
  }
  event index {
    community;
    description;
    type (log | trap | log-and-trap | none);
  }
}

```

If you do not have CLI access, you can configure remote monitoring using the SNMP Manager or management application, assuming SNMP access has been granted. (See [Table 24 on page 259](#).) To configure RMON using SNMP, perform SNMP **Set** requests to the RMON event and alarm tables.

RMON Event Table

Set up an event for each type that you want to generate. For example, you could have two generic events, *rising* and *falling*, or many different events for each variable that is being monitored (for example, *temperature rising* event, *temperature falling* event, *firewall hit* event, *interface utilization* event, and so on). Once the events have been configured, you do not need to update them.

Table 24: RMON Event Table

Field	Description
eventDescription	Text description of this event
eventType	Type of event (for example, log , trap , or log and trap)
eventCommunity	Trap group to which to send this event (as defined in the Junos OS configuration, which is not the same as the community)
eventOwner	Entity (for example, manager) that created this event
eventStatus	Status of this row (for example, valid , invalid , or createRequest)

RMON Alarm Table

The RMON alarm table stores the SNMP object identifiers (including their instances) of the variables that are being monitored, together with any rising and falling thresholds and their corresponding event indexes. To create an RMON request, specify the fields shown in [Table 25 on page 259](#).

Table 25: RMON Alarm Table

Field	Description
alarmStatus	Status of this row (for example, valid , invalid , or createRequest)

Table 25: RMON Alarm Table (*continued*)

Field	Description
alarmInterval	Sampling period (in seconds) of the monitored variable
alarmVariable	OID (and instance) of the variable to be monitored
alarmValue	Actual value of the sampled variable
alarmSampleType	Sample type (absolute or delta changes)
alarmStartupAlarm	Initial alarm (rising , falling , or either)
alarmRisingThreshold	Rising threshold against which to compare the value
alarmFallingThreshold	Falling threshold against which to compare the value
alarmRisingEventIndex	Index (row) of the rising event in the event table
alarmFallingEventIndex	Index (row) of the falling event in the event table

Both the **alarmStatus** and **eventStatus** fields are **entryStatus** primitives, as defined in RFC 2579, *Textual Conventions for SMIV2*.

Troubleshooting RMON

You troubleshoot the RMON agent, **rmopd**, that runs on the router by inspecting the contents of the Juniper Networks enterprise RMON MIB, **jnxRmon**, which provides the extensions listed in [Table 26 on page 260](#) to the RFC 2819 **alarmTable**.

Table 26: jnxRmon Alarm Extensions

Field	Description
jnxRmonAlarmGetFailCnt	Number of times the internal Get request for the variable failed
jnxRmonAlarmGetFailTime	Value of sysUpTime when the last failure occurred
jnxRmonAlarmGetFailReason	Reason why the Get request failed
jnxRmonAlarmGetOkTime	Value of sysUpTime when the variable moved out of failure state
jnxRmonAlarmState	Status of this alarm entry

Monitoring the extensions in this table provides clues as to why remote alarms may not behave as expected.

Related Documentation

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 261](#)

Understanding Measurement Points, Key Performance Indicators, and Baseline Values

Supported Platforms M Series, MX Series, PTX Series, SRX Series, T Series

This chapter topic provides guidelines for monitoring the service quality of an IP network. It describes how service providers and network administrators can use information provided by Juniper Networks routers to monitor network performance and capacity. You should have a thorough understanding of the SNMP and the associated MIB supported by Junos OS.



NOTE: For a good introduction to the process of monitoring an IP network, see RFC 2330, *Framework for IP Performance Metrics*.

This topic contains the following sections:

- [Measurement Points on page 261](#)
- [Basic Key Performance Indicators on page 262](#)
- [Setting Baselines on page 262](#)


Measurement Points

Defining the measurement points where metrics are measured is equally as important as defining the metrics themselves. This section describes measurement points within the context of this chapter and helps identify where measurements can be taken from a service provider network. It is important to understand exactly where a measurement point is. Measurement points are vital to understanding the implication of what the actual measurement means.

An IP network consists of a collection of routers connected by physical links that are all running the Internet Protocol. You can view the network as a collection of routers with an ingress (entry) point and an egress (exit) point. See [Figure 4 on page 261](#).

- Network-centric measurements are taken at measurement points that most closely map to the ingress and egress points for the network itself. For example, to measure delay across the provider network from Site A to Site B, the measurement points should be the ingress point to the provider network at Site A and the egress point at Site B.
- Router-centric measurements are taken directly from the routers themselves, but be careful to ensure that the correct router subcomponents have been identified in advance.

Figure 4: Network Entry Points



g017042



NOTE: [Figure 4 on page 261](#) does not show the client networks at customer premises, but they would be located on either side of the ingress and egress points. Although this chapter does not discuss how to measure network services as perceived by these client networks, you can use measurements taken for the service provider network as input into such calculations.

Basic Key Performance Indicators

For example, you could monitor a service provider network for three basic key performance indicators (KPIs):

- *Availability* measures the “reachability” of one measurement point from another measurement point at the network layer (for example, using ICMP ping). The underlying routing and transport infrastructure of the provider network will support the availability measurements, with failures highlighted as unavailability.
- *Health* measures the number and type of errors that are occurring on the provider network, and can consist of both router-centric and network-centric measurements, such as hardware failures or packet loss.
- *Performance* of the provider network measures how well it can support IP services (for example, in terms of delay or utilization).

Setting Baselines

How well is the provider network performing? We recommend an initial three-month period of monitoring to identify a network’s normal operational parameters. With this information, you can recognize exceptions and identify abnormal behavior. You should continue baseline monitoring for the lifetime of each measured metric. Over time, you must be able to recognize performance trends and growth patterns.

Within the context of this chapter, many of the metrics identified do not have an allowable operational range associated with them. In most cases, you cannot identify the allowable operational range until you have determined a baseline for the actual variable on a specific network.

Related Documentation

- [Understanding RMON for Monitoring Service Quality on page 257](#)
- [Defining and Measuring Network Availability on page 262](#)
- [Measuring Health on page 268](#)
- [Measuring Performance on page 274](#)

Defining and Measuring Network Availability

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, T Series](#)

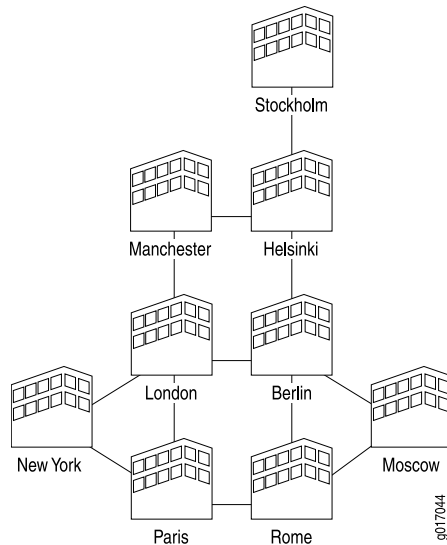
This topic includes the following sections:

- [Defining Network Availability on page 263](#)
- [Measuring Availability on page 265](#)

Defining Network Availability

Availability of a service provider's IP network can be thought of as the reachability between the regional points of presence (POP), as shown in [Figure 5 on page 263](#).


Figure 5: Regional Points of Presence



With the example above, when you use a full mesh of measurement points, where every POP measures the availability to every other POP, you can calculate the total availability of the service provider's network. This KPI can also be used to help monitor the service level of the network, and can be used by the service provider and its customers to determine if they are operating within the terms of their service-level agreement (SLA).

Where a POP may consist of multiple routers, take measurements to each router as shown in [Figure 6 on page 263](#).

Figure 6: Measurements to Each Router



Measurements include:

- Path availability—Availability of an egress interface **B1** as seen from an ingress interface **A1**.
- Router availability—Percentage of path availability of all measured paths terminating on the router.
- POP availability—Percentage of router availability between any two regional POPs, **A** and **B**.
- Network availability—Percentage of POP availability for all regional POPs in the service provider's network.

To measure POP availability of **POP A** to **POP B** in [Figure 6 on page 263](#), you must measure the following four paths:

Path A1 => B1
Path A1 => B2
Path A2 => B1
Path A2 => B2

Measuring availability from **POP B** to **POP A** would require a further four measurements, and so on.

A full mesh of availability measurements can generate significant management traffic. From the sample diagram above:

- Each POP has two co-located provider edge (PE) routers, each with 2xSTM1 interfaces, for a total of 18 PE routers and 36xSTM1 interfaces.
- There are six core provider (P) routers, four with 2xSTM4 and 3xSTM1 interfaces each, and two with 3xSTM4 and 3xSTM1 interfaces each.

This makes a total of 68 interfaces. A full mesh of paths between every interface is:

$[n \times (n-1)] / 2$ gives $[68 \times (68-1)] / 2 = 2278$ paths

To reduce management traffic on the service provider's network, instead of generating a full mesh of interface availability tests (for example, from each interface to every other interface), you can measure from each router's loopback address. This reduces the number of availability measurements required to a total of one for each router, or:

$[n \times (n-1)] / 2$ gives $[24 \times (24-1)] / 2 = 276$ measurements

This measures availability from each router to every other router.

Monitoring the SLA and the Required Bandwidth

A typical SLA between a service provider and a customer might state:

A Point of Presence is the connection of two back-to-back provider edge routers to separate core provider routers using different links for resilience. The system is considered to be unavailable when either an entire POP becomes unavailable or for the duration of a Priority 1 fault.

An SLA availability figure of 99.999 percent for a provider's network would relate to a down time of approximately 5 minutes per year. Therefore, to measure this proactively,

you would have to take availability measurements at a granularity of less than one every five minutes. With a standard size of 64 bytes per ICMP ping request, one ping test per minute would generate 7680 bytes of traffic per hour per destination, including ping responses. A full mesh of ping tests to 276 destinations would generate 2,119,680 bytes per hour, which represents the following:

- On an OC3/STM1 link of 155.52 Mbps, a utilization of 1.362 percent
- On an OC12/STM4 link of 622.08 Mbps, a utilization of 0.340 percent

With a size of 1500 bytes per ICMP ping request, one ping test per minute would generate 180,000 bytes per hour per destination, including ping responses. A full mesh of ping tests to 276 destinations would generate 49,680,000 bytes per hour, which represents the following:

- On an OC3/STM1 link, 31.94 percent utilization
- On an OC12/STM4 link, 7.986 percent utilization

Each router can record the results for every destination tested. With one test per minute to each destination, a total of $1 \times 60 \times 24 \times 276 = 397,440$ tests per day would be performed and recorded by each router. All ping results are stored in the **pingProbeHistoryTable** (see RFC 2925) and can be retrieved by an SNMP performance reporting application (for example, service performance management software from InfoVista, Inc., or Concord Communications, Inc.) for post processing. This table has a maximum size of 4,294,967,295 rows, which is more than adequate.

Measuring Availability

There are two methods you can use to measure availability:

- Proactive—Availability is automatically measured as often as possible by an operational support system.
- Reactive—Availability is recorded by a Help desk when a fault is first reported by a user or a fault monitoring system.

This section discusses real-time performance monitoring as a proactive monitoring solution.

Real-Time Performance Monitoring

Juniper Networks provides a real-time performance monitoring (RPM) service to monitor real-time network performance. Use the J-Web Quick Configuration feature to configure real-time performance monitoring parameters used in real-time performance monitoring tests. (J-Web Quick Configuration is a browser-based GUI that runs on Juniper Networks routers. For more information, see the *J-Web Interface User Guide*.)

Configuring Real-Time Performance Monitoring

Some of the most common options you can configure for real-time performance monitoring tests are shown in [Table 27 on page 266](#).

Table 27: Real-Time Performance Monitoring Configuration Options

Field	Description
Request Information	
Probe Type	Type of probe to send as part of the test. Probe types can be: <ul style="list-style-type: none"> • <code>http-get</code> • <code>http-get-metadata</code> • <code>icmp-ping</code> • <code>icmp-ping-timestamp</code> • <code>tcp-ping</code> • <code>udp-ping</code>
Interval	Wait time (in seconds) between each probe transmission. The range is 1 to 255 seconds.
Test Interval	Wait time (in seconds) between tests. The range is 0 to 86400 seconds.
Probe Count	Total number of probes sent for each test. The range is 1 to 15 probes.
Destination Port	TCP or UDP port to which probes are sent. Use number 7—a standard TCP or UDP port number—or select a port number from 49152 through 65535.
DSCP Bits	Differentiated Services code point (DSCP) bits. This value must be a valid 6-bit pattern. The default is 000000.
Data Size	Size (in bytes) of the data portion of the ICMP probes. The range is 0 to 65507 bytes.
Data Fill	Contents of the data portion of the ICMP probes. Contents must be a hexadecimal value. The range is 1 to 800h.
Maximum Probe Thresholds	
Successive Lost Probes	Total number of probes that must be lost successively to trigger a probe failure and generate a system log message. The range is 0 to 15 probes.
Lost Probes	Total number of probes that must be lost to trigger a probe failure and generate a system log message. The range is 0 to 15 probes.
Round Trip Time	Total round-trip time (in microseconds) from the Services Router to the remote server, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Jitter	Total jitter (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.

Table 27: Real-Time Performance Monitoring Configuration Options (*continued*)

Field	Description
Standard Deviation	Maximum allowable standard deviation (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Egress Time	Total one-way time (in microseconds) from the router to the remote server, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Ingress Time	Total one-way time (in microseconds) from the remote server to the router, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Jitter Egress Time	Total outbound-time jitter (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Jitter Ingress Time	Total inbound-time jitter (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Egress Standard Deviation	Maximum allowable standard deviation of outbound times (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.
Ingress Standard Deviation	Maximum allowable standard deviation of inbound times (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message. The range is 0 to 60,000,000 microseconds.

Displaying Real-Time Performance Monitoring Information

For each real-time performance monitoring test configured on the router, monitoring information includes the round-trip time, jitter, and standard deviation. To view this information, select **Monitor > RPM** in the J-Web interface, or enter the **show services rpm** command-line interface (CLI) command.

To display the results of the most recent real-time performance monitoring probes, enter the **show services rpm probe-results** CLI command:

```
user@host> show services rpm probe-results
Owner: p1, Test: t1
Target address: 10.8.4.1, Source address: 10.8.4.2, Probe type: icmp-ping
Destination interface name: lt-0/0/0.0
Test size: 10 probes
Probe results:
  Response received, Sun Jul 10 19:07:34 2005
  Rtt: 50302 usec
```

Results over current test:

Probes sent: 2, Probes received: 1, Loss percentage: 50

Measurement: Round trip time

Minimum: 50302 usec, Maximum: 50302 usec, Average: 50302 usec,

Jitter: 0 usec, Stddev: 0 usec

Results over all tests:

Probes sent: 2, Probes received: 1, Loss percentage: 50

Measurement: Round trip time

Minimum: 50302 usec, Maximum: 50302 usec, Average: 50302 usec,

Jitter: 0 usec, Stddev: 0 usec

Related Documentation

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 261](#)
- [Understanding RMON for Monitoring Service Quality on page 257](#)
- [Measuring Health on page 268](#)
- [Measuring Performance on page 274](#)

Measuring Health

Supported Platforms [M Series, MX Series, T Series](#)

You can monitor health metrics reactively by using fault management software such as SMARTS InCharge, Micromuse Netcool Omnibus, or Concord Live Exceptions. We recommend that you monitor the health metrics shown in [Table 28 on page 268](#).

Table 28: Health Metrics

Metric:	Errors in
Description	Number of inbound packets that contained errors, preventing them from being delivered
MIB name	IF-MIB (RFC 2233)
Variable name	ifInErrors
Variable OID	.1.3.6.1.31.2.2.1.14
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	Logical interfaces
Metric:	Errors out
Description	Number of outbound packets that contained errors, preventing them from being transmitted
MIB name	IF-MIB (RFC 2233)

Table 28: Health Metrics (*continued*)

Variable name	ifOutErrors
Variable OID	.1.3.6.1.31.2.2.1.20
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	Logical interfaces
Metric:	Discards in
Description	Number of inbound packets discarded, even though no errors were detected
MIB name	IF-MIB (RFC 2233)
Variable name	ifInDiscards
Variable OID	.1.3.6.1.31.2.2.1.13
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	Logical interfaces
Metric:	Unknown protocols
Description	Number of inbound packets discarded because they were of an unknown protocol
MIB name	IF-MIB (RFC 2233)
Variable name	ifInUnknownProtos
Variable OID	.1.3.6.1.31.2.2.1.15
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	Logical interfaces
Metric:	Interface operating status
Description	Operational status of an interface
MIB name	IF-MIB (RFC 2233)

Table 28: Health Metrics (*continued*)

Variable name	ifOperStatus
Variable OID	.1.3.6.1.31.2.2.1.8
Frequency (mins)	15
Allowable range	1 (up)
Managed objects	Logical interfaces
Metric:	Label Switched Path (LSP) state
Description	Operational state of an MPLS label-switched path
MIB name	MPLS-MIB
Variable name	mplsLspState
Variable OID	mplsLspEntry.2
Frequency (mins)	60
Allowable range	2 (up)
Managed objects	All label-switched paths in the network
Metric:	Component operating status
Description	Operational status of a router hardware component
MIB name	JUNIPER-MIB
Variable name	jnxOperatingState
Variable OID	.1.3.6.1.4.1.2636.1.13.1.6
Frequency (mins)	60
Allowable range	2 (running) or 3 (ready)
Managed objects	All components in each Juniper Networks router
Metric:	Component operating temperature
Description	Operational temperature of a hardware component, in Celsius
MIB name	JUNIPER-MIB
Variable name	jnxOperatingTemp

Table 28: Health Metrics (*continued*)

Variable OID	.1.3.6.1.4.1.2636.1.13.1.7
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All components in a chassis
Metric:	System up time
Description	Time, in milliseconds, that the system has been operational.
MIB name	MIB-2 (RFC 1213)
Variable name	sysUpTime
Variable OID	.1.3.6.1.1.3
Frequency (mins)	60
Allowable range	Increasing only (decrement indicates a restart)
Managed objects	All routers
Metric:	No IP route errors
Description	Number of packets that could not be delivered because there was no IP route to their destination.
MIB name	MIB-2 (RFC 1213)
Variable name	ipOutNoRoutes
Variable OID	ip.12
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	Each router
Metric:	Wrong SNMP community names
Description	Number of incorrect SNMP community names received
MIB name	MIB-2 (RFC 1213)
Variable name	snmplnBadCommunityNames

Table 28: Health Metrics (*continued*)

Variable OID	snmp.4
Frequency (hours)	24
Allowable range	To be baselined
Managed objects	Each router
Metric:	SNMP community violations
Description	Number of valid SNMP communities used to attempt invalid operations (for example, attempting to perform SNMP Set requests)
MIB name	MIB-2 (RFC 1213)
Variable name	snmpInBadCommunityUses
Variable OID	snmp.5
Frequency (hours)	24
Allowable range	To be baselined
Managed objects	Each router
Metric:	Redundancy switchover
Description	Total number of redundancy switchovers reported by this entity
MIB name	JUNIPER-MIB
Variable name	jnxRedundancySwitchoverCount
Variable OID	jnxRedundancyEntry.8
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All Juniper Networks routers with redundant Routing Engines
Metric:	FRU state
Description	Operational status of each field-replaceable unit (FRU)
MIB name	JUNIPER-MIB
Variable name	jnxFruState

Table 28: Health Metrics (*continued*)

Variable OID	jnxFruEntry.8
Frequency (mins)	15
Allowable range	2 through 6 for ready/online states. See jnxFruOfflineReason in the event of a FRU failure.
Managed objects	All FRUs in all Juniper Networks routers.
Metric:	Rate of tail-dropped packets
Description	Rate of tail-dropped packets per output queue, per forwarding class, per interface.
MIB name	JUNIPER-COS-MIB
Variable name	jnxCosIfqTailDropPktRate
Variable OID	jnxCosIfqStatsEntry.12
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	For each forwarding class per interface in the provider network, when CoS is enabled.
Metric:	Interface utilization: octets received
Description	Total number of octets received on the interface, including framing characters.
MIB name	IF-MIB
Variable name	ifInOctets
Variable OID	.1.3.6.1.2.1.2.2.1.10.x
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All operational interfaces in the network
Metric:	Interface utilization: octets transmitted
Description	Total number of octets transmitted out of the interface, including framing characters.
MIB name	IF-MIB

Table 28: Health Metrics (*continued*)

Variable name	ifOutOctets
Variable OID	.1.3.6.1.2.1.2.2.1.16.x
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All operational interfaces in the network



NOTE: Byte counts vary depending on interface type, encapsulation used and PIC supported. For example, with vlan-ccc encapsulation on a 4xFE, GE, or GE IQ PIC, the byte count includes framing and control word overhead. (See [Table 29 on page 274](#).)

Table 29: Counter Values for vlan-ccc Encapsulation

PIC Type	Encapsulation	input (Unit Level)	Output (Unit Level)	SNMP
4xFE	vlan-ccc	Frame (no frame check sequence [FCS])	Frame (including FCS and control word)	ifInOctets, ifOutOctets
GE	vlan-ccc	Frame (no FCS)	Frame (including FCS and control word)	ifInOctets, ifOutOctets
GE IQ	vlan-ccc	Frame (no FCS)	Frame (including FCS and control word)	ifInOctets, ifOutOctets

SNMP traps are also a good mechanism to use for health management. For more information, see [“Standard SNMP Traps Supported by Junos OS” on page 50](#) and [“Enterprise-Specific SNMP Traps Supported by Junos OS” on page 57](#).

Related Documentation

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 261](#)
- [Understanding RMON for Monitoring Service Quality on page 257](#)
- [Defining and Measuring Network Availability on page 262](#)
- [Measuring Performance on page 274](#)
- [SNMB MIB Explorer](#)

Measuring Performance

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

The performance of a service provider's network is usually defined as how well it can support services, and is measured with metrics such as delay and utilization. We suggest that you monitor the following performance metrics using applications such as InfoVista Service Performance Management or Concord Network Health (see [Table 30 on page 275](#)).

Table 30: Performance Metrics

Metric:	Average delay
Description	Average round-trip time (in milliseconds) between two measurement points.
MIB name	DISMAN-PING-MIB (RFC 2925)
Variable name	pingResultsAverageRtt
Variable OID	pingResultsEntry.6
Frequency (mins)	15 (or depending upon ping test frequency)
Allowable range	To be baselined
Managed objects	Each measured path in the network
Metric:	Interface utilization
Description	Utilization percentage of a logical connection.
MIB name	IF-MIB
Variable name	(ifInOctets & ifOutOctets) * 8 / ifSpeed
Variable OID	ifTable entries
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All operational interfaces in the network
Metric:	Disk utilization
Description	Utilization of disk space within the Juniper Networks router
MIB name	HOST-RESOURCES-MIB (RFC 2790)
Variable name	hrStorageSize – hrStorageUsed
Variable OID	hrStorageEntry.5 – hrStorageEntry.6
Frequency (mins)	1440

Table 30: Performance Metrics (*continued*)

Allowable range	To be baselined
Managed objects	All Routing Engine hard disks
Metric:	Memory utilization
Description	Utilization of memory on the Routing Engine and FPC.
MIB name	JUNIPER-MIB (Juniper Networks enterprise Chassis MIB)
Variable name	jnxOperatingHeap
Variable OID	Table for each component
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All Juniper Networks routers
Metric:	CPU load
Description	Average utilization over the past minute of a CPU.
MIB name	JUNIPER-MIB (Juniper Networks enterprise Chassis MIB)
Variable name	jnxOperatingCPU
Variable OID	Table for each component
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	All Juniper Networks routers
Metric:	LSP utilization
Description	Utilization of the MPLS label-switched path.
MIB name	MPLS-MIB
Variable name	mplsPathBandwidth / (mplsLspOctets * 8)
Variable OID	mplsLspEntry.21 and mplsLspEntry.3
Frequency (mins)	60
Allowable range	To be baselined

Table 30: Performance Metrics (*continued*)

Managed objects	All label-switched paths in the network
Metric:	Output queue size
Description	Size, in packets, of each output queue per forwarding class, per interface.
MIB name	JUNIPER-COS-MIB
Variable name	jnxCosIfqQedPkts
Variable OID	jnxCosIfqStatsEntry.3
Frequency (mins)	60
Allowable range	To be baselined
Managed objects	For each forwarding class per interface in the network, once CoS is enabled.

This section includes the following topics:

- [Measuring Class of Service on page 277](#)
- [Inbound Firewall Filter Counters per Class on page 278](#)
- [Monitoring Output Bytes per Queue on page 279](#)
- [Dropped Traffic on page 280](#)


Measuring Class of Service

You can use class-of-service (CoS) mechanisms to regulate how certain classes of packets are handled within your network during times of peak congestion. Typically you must perform the following steps when implementing a CoS mechanism:

- Identify the type of packets that is applied to this class. For example, include all customer traffic from a specific ingress edge interface within one class, or include all packets of a particular protocol such as voice over IP (VoIP).
- Identify the required deterministic behavior for each class. For example, if VoIP is important, give VoIP traffic the highest priority during times of network congestion. Conversely, you can downgrade the importance of Web traffic during congestion, as it may not impact customers too much.

With this information, you can configure mechanisms at the network ingress to monitor, mark, and police traffic classes. Marked traffic can then be handled in a more deterministic way at egress interfaces, typically by applying different queuing mechanisms for each class during times of network congestion. You can collect information from the network to provide customers with reports showing how the network is behaving during times of congestion. (See [Figure 7 on page 278](#).)

Figure 7: Network Behavior During Congestion



To generate these reports, routers must provide the following information:

- Submitted traffic—Amount of traffic received per class.
- Delivered traffic—Amount of traffic transmitted per class.
- Dropped traffic—Amount of traffic dropped because of CoS limits.

The following section outlines how this information is provided by Juniper Networks routers.

Inbound Firewall Filter Counters per Class

Firewall filter counters are a very flexible mechanism you can use to match and count inbound traffic per class, per interface. For example:

```
firewall {
  filter f1 {
    term t1 {
      from {
        dscp af11;
      }
      then {
        # Assured forwarding class 1 drop profile 1 count inbound-af11;
        accept;
      }
    }
  }
}
```

For example, [Table 31 on page 278](#) shows additional filters used to match the other classes.

Table 31: Inbound Traffic Per Class

DSCP Value	Firewall Match Condition	Description
10	af11	Assured forwarding class 1 drop profile 1
12	af12	Assured forwarding class 1 drop profile 2

Table 31: Inbound Traffic Per Class (continued)

DSCP Value	Firewall Match Condition	Description
18	af21	Best effort class 2 drop profile 1
20	af22	Best effort class 2 drop profile 2
26	af31	Best effort class 3 drop profile 1

Any packet with a CoS DiffServ code point (DSCP) conforming to RFC 2474 can be counted in this way. The Juniper Networks enterprise-specific Firewall Filter MIB presents the counter information in the variables shown in [Table 32 on page 279](#).

Table 32: Inbound Counters

Indicator Name	Inbound Counters
MIB	jnxFirewalls
Table	jnxFirewallCounterTable
Index	jnxFWFilter.jnxFWCounter
Variables	jnxFWCounterPacketCount jnxFWCounterByteCount
Description	Number of bytes being counted pertaining to the specified firewall filter counter
SNMP version	SNMPv2

This information can be collected by any SNMP management application that supports SNMPv2. Products from vendors such as Concord Communications, Inc., and InfoVista, Inc., provide support for the Juniper Networks Firewall MIB with their native Juniper Networks device drivers.

Monitoring Output Bytes per Queue

You can use the Juniper Networks enterprise ATM CoS MIB to monitor outbound traffic, per virtual circuit forwarding class, per interface. (See [Table 33 on page 279](#).)

Table 33: Outbound Counters for ATM Interfaces

Indicator Name	Outbound Counters
MIB	JUNIPER-ATM-COS-MIB
Variable	jnxCosAtmVcQstatsOutBytes
Index	ifIndex.atmVclVpi.atmVclVci.jnxCosFclId

Table 33: Outbound Counters for ATM Interfaces (*continued*)

Indicator Name	Outbound Counters
Description	Number of bytes belonging to the specified forwarding class that were transmitted on the specified virtual circuit.
SNMP version	SNMPv2

Non-ATM interface counters are provided by the Juniper Networks enterprise-specific CoS MIB, which provides information shown in [Table 34 on page 280](#).

Table 34: Outbound Counters for Non-ATM Interfaces

Indicator Name	Outbound Counters
MIB	JUNIPER-COS-MIB
Table	jnxCosIfqStatsTable
Index	jnxCosIfqIfIndex.jnxCosIfqFc
Variables	jnxCosIfqTxedBytes jnxCosIfqTxedPkts
Description	Number of transmitted bytes or packets per interface per forwarding class
SNMP version	SNMPv2

Dropped Traffic

You can calculate the amount of dropped traffic by subtracting the outbound traffic from the incoming traffic:

$$\text{Dropped} = \text{Inbound Counter} - \text{Outbound Counter}$$

You can also select counters from the CoS MIB, as shown in [Table 35 on page 280](#).

Table 35: Dropped Traffic Counters

Indicator Name	Dropped Traffic
MIB	JUNIPER-COS-MIB
Table	jnxCosIfqStatsTable
Index	jnxCosIfqIfIndex.jnxCosIfqFc
Variables	jnxCosIfqTailDropPkts jnxCosIfqTotalRedDropPkts

Table 35: Dropped Traffic Counters (*continued*)

Indicator Name	Dropped Traffic
Description	The number of tail-dropped or RED-dropped packets per interface per forwarding class
SNMP version	SNMPv2

**Related
Documentation**

- [Understanding Measurement Points, Key Performance Indicators, and Baseline Values on page 261](#)
- [Understanding RMON for Monitoring Service Quality on page 257](#)
- [Defining and Measuring Network Availability on page 262](#)
- [Measuring Health on page 268](#)

PART 4

Health Monitoring with SNMP

- [Configuring Health Monitoring on page 285](#)

Configuring Health Monitoring

- [Configuring Health Monitoring on Devices Running Junos OS on page 285](#)
- [Example: Configuring Health Monitoring on page 288](#)

Configuring Health Monitoring on Devices Running Junos OS

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

As the number of devices managed by a typical network management system (NMS) grows and the complexity of the devices themselves increases, it becomes increasingly impractical for the NMS to use polling to monitor the devices. A more scalable approach is to rely on network devices to notify the NMS when something requires attention.

On Juniper Networks routers, RMON alarms and events provide much of the infrastructure needed to reduce the polling overhead from the NMS. However, with this approach, you must set up the NMS to configure specific MIB objects into RMON alarms. This often requires device-specific expertise and customizing of the monitoring application. In addition, some MIB object instances that need monitoring are set only at initialization or change at runtime and cannot be configured in advance.

To address these issues, the health monitor extends the RMON alarm infrastructure to provide predefined monitoring for a selected set of object instances (for file system usage, CPU usage, and memory usage) and includes support for unknown or dynamic object instances (such as Junos OS processes).

Health monitoring is designed to minimize user configuration requirements. To configure health monitoring entries, include the **health-monitor** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
health-monitor {
  falling-threshold percentage;
  interval seconds;
  rising-threshold percentage;
  idp {
    falling-threshold percentage;
    interval seconds;
    rising-threshold percentage;
  }
}
```

Configuring monitoring events at the **[edit snmp health-monitor]** hierarchy level sets polling intervals for the overall system health. If you set these same options at the **[edit snmp health-monitor idp]** hierarchy level, an SNMP event is generated by the device if the percentage of dataplane memory utilized by the intrusion detection and prevention (IDP) system rises above or falls below your settings.

You can use the **show snmp health-monitor** operational command to view information about health monitor alarms and logs.

This topic describes the minimum required configuration and discusses the following tasks for configuring the health monitor:

- [Monitored Objects on page 286](#)
- [Minimum Health Monitoring Configuration on page 287](#)
- [Configuring the Falling Threshold or Rising Threshold on page 287](#)
- [Configuring the Interval on page 288](#)
- [Log Entries and Traps on page 288](#)

Monitored Objects

When you configure the health monitor, monitoring information for certain object instances is available, as shown in [Table 36 on page 286](#).

Table 36: Monitored Object Instances

Object	Description
jnxFileUsed jnxFileUsed	Monitors the following file system on the router or switch: /dev/ad0s1a: This is the root file system mounted on /.
jnxFileUsed2 jnxFileUsed2	Monitors the following file system on the router or switch: /dev/ad0s1e: This is the configuration file system mounted on /config.
jnxOperatingCPU jnxOperatingCPU (RE0)	Monitors CPU usage for Routing Engines (RE0 and RE1). The index values assigned to Routing Engines depend on whether the Chassis MIB uses a zero-based or ones-based indexing scheme. Because the indexing scheme is configurable, the proper index is determined when the router or switch is initialized and when there is a configuration change. If the router or switch has only one Routing Engine, the alarm entry monitoring RE1 is removed after five failed attempts to obtain the CPU value.
jnxOperatingCPU jnxOperatingCPU (RE1)	
jnxOperatingBuffer jnxOperatingBuffer (RE0)	Monitors the amount of memory available on Routing Engines (RE0 and RE1). Because the indexing of this object is identical to that used for jnxOperatingCPU , index values are adjusted depending on the indexing scheme used in the Chassis MIB. As with jnxOperatingCPU , the alarm entry monitoring RE1 is removed if the router or switch has only one Routing Engine.
jnxOperatingBuffer jnxOperatingBuffer (RE1)	

Table 36: Monitored Object Instances (*continued*)

Object	Description
snmp AppRtCPU	Monitors the CPU usage for each Junos OS process (also called daemon). Multiple instances of the same process are monitored and indexed separately.
snmp AppRtMemory	Monitors the memory usage for each Junos OS process. Multiple instances of the same process are monitored and indexed separately.

Minimum Health Monitoring Configuration

To enable health monitoring on the router or switch, include the **health-monitor** statement at the **[edit snmp]** hierarchy level:

```
[edit snmp]
health-monitor;
```

Configuring the Falling Threshold or Rising Threshold

The falling threshold is the lower threshold (expressed as a percentage of the maximum possible value) for the monitored variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the rising threshold. You must specify the falling threshold as a percentage of the maximum possible value. The default is **70** percent.

By default, the rising threshold is **80** percent of the maximum possible value for the monitored object instance. The rising threshold is the upper threshold for the monitored variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the falling threshold. You must specify the rising threshold as a percentage of the maximum possible value for the monitored variable.

To configure the falling threshold or rising threshold, include the **falling-threshold** or **rising-threshold** statement at the **[edit snmp health-monitor]** hierarchy level:

```
[edit snmp health-monitor]
falling-threshold percentage;
rising-threshold percentage;
```

percentage can be a value from 1 through 100.

The falling and rising thresholds apply to all object instances monitored by the health monitor.

Configuring the Interval

The interval represents the period of time, in seconds, over which the object instance is sampled and compared with the rising and falling thresholds.

To configure the interval, include the **interval** statement and specify the number of seconds at the **[edit snmp health-monitor]** hierarchy level:

```
[edit snmp health-monitor]
interval seconds;
```

seconds can be a value from 1 through 2147483647. The default is 300 seconds (5 minutes).

Log Entries and Traps

The system log entries generated for any health monitor events (thresholds crossed, errors, and so on) have a corresponding **HEALTHMONITOR** tag rather than a generic **SNMPD_RMON_EVENTLOG** tag. However, the health monitor sends generic RMON **risingThreshold** and **fallingThreshold** traps.

Related Documentation

- [Understanding RMON Alarms and Events Configuration on page 241](#)
- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
- [Configuring an RMON Event Entry and Its Attributes on page 246](#)
- [Example: Configuring Health Monitoring on page 288](#)
- [Understanding Device Management Functions in Junos OS on page 3](#)
- [health-monitor on page 685](#)

Example: Configuring Health Monitoring

Supported Platforms **M Series, MX Series, PTX Series, SRX Series, T Series**

Configure the health monitor:

```
[edit snmp]
health-monitor {
  falling-threshold 85;
  interval 600;
  rising-threshold 75;
}
```

In this example, the sampling interval is every 600 seconds (10 minutes), the falling threshold is 85 percent of the maximum possible value for each object instance monitored, and the rising threshold is 75 percent of the maximum possible value for each object instance monitored.

Related Documentation

- [Configuring Health Monitoring on Devices Running Junos OS on page 285](#)

PART 5

Gathering Statistics for Accounting Purposes Using Accounting Options, Source Class Usage and Destination Class Usage Options

- [Accounting Options, Source Class Usage and Destination Class Usage Options Overview on page 291](#)
- [Configuring Accounting Options, Source Class Usage and Destination Class Usage Options on page 295](#)

Accounting Options, Source Class Usage and Destination Class Usage Options Overview

- [Accounting Options Overview on page 291](#)
- [Understanding Source Class Usage and Destination Class Usage Options on page 292](#)

Accounting Options Overview

Supported Platforms [ACX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX](#)

An accounting profile represents common characteristics of collected accounting data, including the following:

- Collection interval
- File to contain accounting data
- Specific fields and counter names on which to collect statistics

You can configure multiple accounting profiles, as described in [Table 37 on page 291](#).

Table 37: Types of Accounting Profiles

Type of Profile	Description
Interface profile	Collects the specified error and statistic information.
Filter profile	Collects the byte and packet counts for the counter names specified in the filter profile.
MIB profile	Collects selected MIB statistics and logs them to a specified file.
Routing Engine profile	Collects selected Routing Engine statistics and logs them to a specified file.
Class usage profile	Collects class usage statistics and logs them to a specified file.

Related Documentation

- [Understanding Device Management Functions in Junos OS on page 3](#)
- [Accounting Options Configuration](#)
- [Configuring Accounting-Data Log Files on page 304](#)
- [Configuring the Interface Profile](#)
- [Configuring the Filter Profile on page 312](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level on page 295](#)

Understanding Source Class Usage and Destination Class Usage Options

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

You can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets defined as source classes and *destination classes*. You can define classes based on a variety of parameters, such as routing neighbors, autonomous systems, and route filters.

Source class usage (SCU) counts packets sent to customers by performing lookups on the IP source address and the IP destination address. SCU makes it possible to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. You must enable SCU accounting on both the inbound and outbound physical interfaces.

Destination class usage (DCU) counts packets from customers by performing lookups of the IP destination address. DCU makes it possible to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.

On T Series Core Routers and M320 Multiservice Edge Routers, the source class and destination classes are not carried across the platform fabric. The implications of this are as follows:

- On T Series and M320 routers, SCU and DCU accounting is performed before the packet enters the fabric.
- On T Series and M320 routers, DCU is performed before output filters are evaluated.
- On M Series platforms, DCU is performed after output filters are evaluated.
- If an output filter drops traffic on M Series devices, the dropped packets are excluded from DCU statistics.
- If an output filter drops traffic on T Series and M320 routers, the dropped packets are included in DCU statistics.



NOTE: SCU and DCU is supported on PTX series routers when *enhanced-mode* is configured on the chassis.

On MX Series platforms with MPC/MIC interfaces, SCU and DCU are performed after output filters are evaluated. Packets dropped by output filters are not included in SCU or DCU statistics.

On MX Series platforms with non-MPC/MIC interfaces, SCU and DCU are performed before output filters are evaluated. Packets dropped by output filters are included in SCU and DCU statistics.

On PTX Series platforms, SCU and DCU accounting is performed before output filters are evaluated. Packets dropped by output filters are included in SCU and DCU statistics.

On Enhanced Scaling FPCs (T640-FPC1-ES, T640-FPC2-ES, T640-FPC3-ES, T640-FPC4-1P-ES, and T1600-FPC4-ES), the source class accounting is performed at ingress. Starting with Junos OS Release 14.2, the SCU accounting is performed at ingress on a T4000 Type 5 FPC. The implications of this are as follows:

- SCU accounting is performed when packets traverse from T4000 Type 5 FPC (ingress FPC) to Enhanced Scaling FPCs (egress FPC).
- SCU accounting is performed when packets traverse from Enhanced Scaling FPCs (ingress FPC) to T4000 Type 5 FPC (egress FPC).



NOTE: When the interface statistics are cleared and then the routing engine is replaced, the SCU and DCU statistics will not match the statistics of the previous routing engine.

For more information about source class usage, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide* and the *Junos OS Network Interfaces Library for Routing Devices*.

Related Documentation

- [Example: Grouping Source and Destination Prefixes into a Forwarding Class](#)
- [Configuring SCU or DCU on page 316](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 318](#)
- [Configuring Class Usage Profiles on page 320](#)
- [Configuring the MIB Profile on page 322](#)
- [Configuring the Routing Engine Profile on page 325](#)

CHAPTER 18

Configuring Accounting Options, Source Class Usage and Destination Class Usage Options

- Configuration Statements at the [edit accounting-options] Hierarchy Level on page 295
- Accounting Options Configuration on page 296
- Configuring Accounting-Data Log Files on page 304
- Configuring the Interface Profile on page 309
- Configuring the Filter Profile on page 312
- Example: Configuring a Filter Profile on page 314
- Example: Configuring Interface-Specific Firewall Counters and Filter Profiles on page 315
- Configuring SCU or DCU on page 316
- Configuring SCU on a Virtual Loopback Tunnel Interface on page 318
- Configuring Class Usage Profiles on page 320
- Configuring the MIB Profile on page 322
- Configuring the Routing Engine Profile on page 325

Configuration Statements at the [edit accounting-options] Hierarchy Level

Supported Platforms M Series, MX Series, SRX Series, T Series

This topic shows all possible configuration statements at the [edit accounting-options] hierarchy level and their level in the configuration hierarchy. When you are configuring Junos OS, your current hierarchy level is shown in the banner on the line preceding the **user@host#** prompt.

```
[edit]
accounting-options {
  class-usage-profile profile-name {
    file filename;
    interval minutes;
    destination-classes {
      destination-class-name;
    }
    source-classes {
```

```
        source-class-name;
    }
}
file filename {
    archive-sites {
    }
    files number;
    nonpersistent;
    size bytes;
    start-time time;
    transfer-interval minutes;
}
filter-profile profile-name {
    counters {
        counter-name;
    }
    file filename;
    interval minutes;
}
}
interface-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
mib-profile profile-name {
    file filename;
    interval seconds;
    object-names {
        mib-object-name;
    }
    operation operation-name;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
```

- Related Documentation**
- [Accounting Options Overview on page 291](#)
 - [Accounting Options Configuration on page 296](#)

Accounting Options Configuration

Supported Platforms M Series, MX Series, PTX Series, SRX Series, T Series, vSRX

This topic contains the following sections:

- [Accounting Options—Full Configuration on page 297](#)
- [Minimum Accounting Options Configuration on page 300](#)

Accounting Options—Full Configuration

To configure accounting options, include the following statements at the **[edit accounting-options]** hierarchy level:

```
accounting-options {
  class-usage-profile profile-name {
    file filename;
    interval minutes;
    destination-classes {
      destination-class-name;
    }
    source-classes {
      source-class-name;
    }
    file filename {
      archive-sites {
        site-name;
      }
      files number;
      nonpersistent;
      size bytes;
      source-classes time;
      transfer-interval minutes;
    }
    filter-profile profile-name {
      counters {
        counter-name;
      }
      file filename;
      interval minutes;
    }
  }
  flat-file-profile profile-name {
    fields {
      all-fields;
      egress-stats {
        all-fields;
        input-bytes;
        input-packets;
        output-bytes;
        output-packets;
        queue-id;
        red-drop-bytes;
        red-drop-packets;
        tail-drop-packets;
      }
      general-param {
        all-fields;
        accounting-type;
      }
    }
  }
}
```

```
    descr;  
    line-id;  
    logical-interface;  
    nas-port-id;  
    physical-interface;  
    routing-instance;  
    timestamp;  
    vlan-id;  
  }  
  ingress-stats {  
    all-fields;  
    drop-packets;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
  }  
  l2-stats {  
    all-fields;  
    input-mcast-bytes;  
    input-mcast-packets;  
  }  
  fields {  
    all-fields;  
    egress-stats {  
      all-fields;  
      input-bytes;  
      input-packets;  
      output-bytes;  
      output-packets;  
      queue-id;  
      red-drop-bytes;  
      red-drop-packets;  
      tail-drop-packets;  
    }  
    general-param {  
      all-fields;  
      accounting-type;  
      descr;  
      line-id;  
      logical-interface;  
      nas-port-id;  
      physical-interface;  
      routing-instance;  
      timestamp;  
      vlan-id;  
    }  
    ingress-stats {  
      all-fields;  
      drop-packets;  
      input-bytes;  
      input-packets;  
      output-bytes;  
      output-packets;  
      queue-id;
```

```
}
general-param {
  all-fields;
  accounting-type;
  descr;
  line-id;
  logical-interface;
  nas-port-id;
  physical-interface;
  routing-instance;
  timestamp;
  vlan-id;
}
ingress-stats {
  all-fields;
  drop-packets;
  input-bytes;
  input-packets;
  output-bytes;
  output-packets;
  queue-id;
}
l2-stats {
  all-fields;
  input-mcast-bytes;
  input-mcast-packets;
}
overall-packet {
  all-fields;
  input-bytes;
  input-discards;
  input-errors;
  input-packets;
  inputv6-bytes;
  inputv6-packets;
  output-bytes;
  output-errors;
  output-packets;
  outputv6-bytes;
  outputv6-packets;
}
}
file filename;
format (csv | ipdr)
interval minutes;
schema-version schema-name;
}
interface-profile profile-name {
  fields {
    field-name;
  }
  file filename;
  interval minutes;
}
mib-profile profile-name {
  file filename;
```

```

    interval seconds;
    object-names {
        mib-object-name;
    }
    operation operation-name;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}
}

```

By default, accounting options are disabled.



NOTE: Do not configure MIB objects related to interface octets or packets for a MIB profile, because doing so can cause the SNMP walk or a CLI show command to time out.

Minimum Accounting Options Configuration

To enable accounting options on the router, you must perform at least the following tasks:

- Configure accounting options by including a **file** statement and one or more **source-class-usage**, **destination-class-profile**, **filter-profile**, **interface-profile**, **mib-profile**, or **routing-engine-profile** statements at the **[edit accounting-options]** hierarchy level:

```

[edit]
accounting-options {
    class-usage-profile profile-name {
        file filename;
        interval minutes;
        source-classes {
            source-class-name;
        }
        destination-classes {
            destination-class-name;
        }
    }
    file filename {
        archive-sites {
            site-name;
        }
        files number;
        size bytes;
        transfer-interval minutes;
    }
    filter-profile profile-name {
        counters {
            counter-name;
        }
    }
}

```

```
file filename;
interval minutes;
}
flat-file-profile profile-name{
  fields {
    all-fields;
    egress-stats {
      all-fields;
      input-bytes;
      input-packets;
      output-bytes;
      output-packets;
      queue-id;
      red-drop-bytes;
      red-drop-packets;
      tail-drop-packets;
    }
    general-param {
      all-fields;
      accounting-type;
      descr;
      line-id;
      logical-interface;
      nas-port-id;
      physical-interface;
      routing-instance;
      timestamp;
      vlan-id;
    }
    ingress-stats {
      all-fields;
      drop-packets;
      input-bytes;
      input-packets;
      output-bytes;
      output-packets;
      queue-id;
    }
    l2-stats {
      all-fields;
      input-mcast-bytes;
      input-mcast-packets;
    }
    overall-packet {
      all-fields;
      input-bytes;
      input-discards;
      input-errors;
      input-packets;
      inputv6-bytes;
      inputv6-packets;
      output-bytes;
      output-errors;
      output-packets;
      outputv6-bytes;
      outputv6-packets;
    }
  }
}
```

```
    }  
  }  
  file filename;  
  format (csv | ipdr)  
  interval minutes;  
  schema-version schema-name;  
}  
flat-file-profile profile-name {  
  fields {  
    all-fields;  
    egress-stats {  
      all-fields;  
      input-bytes;  
      input-packets;  
      output-bytes;  
      output-packets;  
      queue-id;  
      red-drop-bytes;  
      red-drop-packets;  
      tail-drop-packets;  
    }  
    general-param {  
      all-fields;  
      accounting-type;  
      descr;  
      line-id;  
      logical-interface;  
      nas-port-id;  
      physical-interface;  
      routing-instance;  
      timestamp;  
      vlan-id;  
    }  
    ingress-stats {  
      all-fields;  
      drop-packets;  
      input-bytes;  
      input-packets;  
      output-bytes;  
      output-packets;  
      queue-id;  
    }  
    l2-stats {  
      all-fields;  
      input-mcast-bytes;  
      input-mcast-packets;  
    }  
    overall-packet {  
      all-fields;  
      input-bytes;  
      input-discards;  
      input-errors;  
      input-packets;  
      inputv6-bytes;  
      inputv6-packets;  
      output-bytes;
```



```

        output-errors;
        output-packets;
        outputv6-bytes;
        outputv6-packets;
    }
}
file filename;
format (csv | ipdr)
interval minutes;
schema-version schema-name;
}
interface-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
mib-profile profile-name {
    file filename;
    interval minutes;
    object-names {
        mib-object-name;
    }
    operation operation-name;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}
}
}

```

- Apply the profiles to the chosen interfaces or filters.

Apply an interface profile to a physical or logical interface by including the **accounting-profile** statement at either the **[edit interfaces *interface-name*]** or the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

```

[edit interfaces]
interface-name {
    accounting-profile profile-name;
    unit logical-unit-number {
        accounting-profile profile-name;
    }
}

```



NOTE: You do not apply destination class profiles to interfaces. Although the interface needs to have the **destination-class-usage** statement configured, the destination class profile automatically finds all interfaces with the destination class configured.

Apply a filter profile to a firewall filter by including the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** hierarchy level:

```
[edit firewall]
filter filter-name {
  accounting-profile profile-name;
}
```

You do not need to apply the Routing Engine profile to an interface because the statistics are collected on the Routing Engine itself.

**Related
Documentation**

- [Accounting Options Overview on page 291](#)
- [Understanding Device Management Functions in Junos OS on page 3](#)
- [Configuring Accounting-Data Log Files on page 304](#)
- [Configuring the Interface Profile on page 309](#)
- [Configuring the Filter Profile on page 312](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level on page 295](#)

Configuring Accounting-Data Log Files

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

An accounting profile specifies what statistics to collect and write to a log file. To configure an accounting-data log file, include the **file** statement at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
cleanup-interval {
  interval days;
}
file filename {
  archive-sites {
    site-name;
  }
  backup-on-failure (master-and-slave | master-only);
  files number;
  nonpersistent;
  push-backup-to-master;
  size bytes;
  start-time time;
  transfer-interval minutes;
}
```

where *filename* is the name of the file in which to write accounting data.

If the filename contains spaces, enclose it in quotation marks (" "). The filename cannot contain a forward slash (/). The file is created in the **/var/log** directory and can contain data from multiple profiles.

All accounting-data log files include header and trailer sections that start with a # in the first column. The header contains the file creation time, the hostname, and the columns that appear in the file. The trailer contains the time that the file was closed.

Whenever any configured value changes that affects the columns in a file, the file creates a new profile layout record that contains a new list of columns.

You must configure the file size; all other properties are optional.

- [Configuring How Long Backup Files Are Retained on page 305](#)
- [Configuring the Maximum Size of the File on page 305](#)
- [Configuring Archive Sites for the Files on page 306](#)
- [Configuring Local Backup for Accounting Files on page 306](#)
- [Configuring Files to Be Compressed on page 307](#)
- [Configuring the Maximum Number of Files on page 307](#)
- [Configuring the Storage Location of the File on page 307](#)
- [Configuring Files to Be Saved After a Change in Mastership on page 308](#)
- [Configuring the Start Time for File Transfer on page 308](#)
- [Configuring the Transfer Interval of the File on page 308](#)

Configuring How Long Backup Files Are Retained

You can configure how many days the files are retained in the local directory before they are deleted.



NOTE: Files saved to the `/var/log/pfedBackup` directory are always compressed to conserve local storage, regardless of whether the `compress` statement is configured.

To configure retention for backup files:

- Specify the number of days.

```
[edit accounting-options]  
user@host# set cleanup-interval interval days
```



NOTE: Files are retained for 1 day if you do not configure this option.

This value, whether configured or default, applies to all configured files at the `[edit accounting-options file]` hierarchy level.

Configuring the Maximum Size of the File

To configure the maximum size of the file:

- Specify the size.

```
[edit accounting-options file filename]
size bytes;
```

The **size** statement is the maximum size of the log file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). The minimum value for **bytes** is 256 KB. You must configure **bytes**; the remaining attributes are optional.

Configuring Archive Sites for the Files

After a file reaches its maximum size or the **transfer-interval** time is exceeded, the file is closed, renamed, and, if you configured an archive site, transferred to a remote host.

To configure the sites where files are archived:

- Specify one or more site names.

```
[edit accounting-options file filename]
user@host# set archive-sites site-name
```

where **site-name** is any valid FTP URL. For more information about specifying valid FTP URLs, see the *Junos OS Administration Library*. You can specify more than one URL, in any order. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, trying the next site in the list only if the transfer does not succeed.

The log file is stored at the archive site with a filename of the format **router-name_log-filename_timestamp**.

Configuring Local Backup for Accounting Files

You can configure the router to save a copy of the accounting file locally when the normal transfer of the files to the archive site fails. The file is saved to the **/var/log/pfedBackup** directory of the relevant Routing Engine. You must specify whether only the files from the master Routing Engine are saved or files are saved from both the master Routing Engine and the backup (slave) Routing Engine.



NOTE: Files saved to the **/var/log/pfedBackup** directory are always compressed to conserve local storage, regardless of whether the **compress** statement is configured.

To configure local backup in the event of failure:

- Specify local backup and which files are saved.

```
[edit accounting-options file filename]
user@host# set backup-on-failure (master-and-slave | master-only)
```

Disabling this feature deletes the backed-up accounting files from the directory.



NOTE: When you do not configure this option, the file is saved on failure into the local directory specified as the last site in the list of archive sites.

Configuring Files to Be Compressed

By default, accounting files are transferred in an uncompressed format. To conserve resources during transmission and on the archive site, you can configure compression for the files.



NOTE: Files saved to the `/var/log/pfedBackup` directory are always compressed to conserve local storage, regardless of whether the `compress` statement is configured.

To configure the router to compress accounting files when they are transferred:

- Specify compression.

```
[edit accounting-options file filename]
user@host# set compress
```

Configuring the Maximum Number of Files

To configure the maximum number of files:

- Specify the number.

```
[edit accounting-options file filename]
user@host# set files number
```

When a log file reaches its maximum size, it is renamed `filename.0`, then `filename.1`, and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. The minimum value for `number` is 3 and the default value is 10.

Configuring the Storage Location of the File

On J Series Services Routers, the files are stored by default on the compact flash drive. Alternatively, you can configure the files to be stored in the `mfs/var/log` directory (on DRAM) instead of the `cf/var/log` directory (on the compact flash drive).

To configure the storage location on DRAM:

- Specify nonpersistent storage.

```
[edit accounting-options file filename]
user@host# set nonpersistent
```

This feature is useful for minimizing read/write traffic on the router's compact flash drive.



NOTE: If log files for accounting data are stored on DRAM, these files are lost when you reboot the router. We recommend that you back up these files periodically.

Configuring Files to Be Saved After a Change in Mastership

You can configure the router to save the accounting files from the new backup Routing Engine to the new master Routing Engine when a change in mastership occurs. The files are stored in the `/var/log/pfedBackup` directory on the router. The master Routing Engine includes these accounting files with its own current accounting files when it transfers the files from the backup directory to the archive site at the next transfer interval. Configure this option when the new backup Routing Engine is not able to connect to the archive site; for example, when the site is not connected by means of an out-of-band interface or the path to the site is routed through a line card.

To configure the backup Routing Engine files to be saved when mastership changes:

- Specify the backup.

```
[edit accounting-options file filename]
user@host# set push-backup-to-master
```



NOTE: The backup Routing Engine's files on the master Routing Engine are sent at each interval even though the files remain the same. If this is more activity than you want, consider using the `backup-on-failure master-and-slave` statement instead.

Configuring the Start Time for File Transfer

To configure the start time for transferring files:

- Specify the time.

```
[edit accounting-options file filename]
user@host# set start-time YYYY-MM-DD.hh:mm
```

For example, 10:00 a.m. on January 30, 2007 is represented as `2007-01-30.10:00`.

Configuring the Transfer Interval of the File

To configure the interval at which files are transferred:

- Specify the interval.

```
[edit accounting-options file filename]
user@host# set transfer-interval minutes
```

The range for `transfer-interval` is 5 through 2880 minutes. The default is 30 minutes.



TIP:

Junos OS saves the existing log file and creates a new file at the configured transfer intervals irrespective of whether:

- The file has reached the maximum size.

- An archive site is configured.

When you have a relatively small transfer interval configured and if no archive site is configured, data can be lost as Junos OS overwrites the log files when the maximum number of log files is reached. To ensure that the log information is saved for a reasonably long time:

- Configure an archive site to archive the log files every time a new log file is created.
 - Configure the maximum value (2880 minutes) for transfer-interval so that new files are created less frequently; that is, only when the file exceeds the maximum size limit or once in 2 days.
-

**Related
Documentation**

- [Accounting Options Overview on page 291](#)
- [Understanding Device Management Functions in Junos OS on page 3](#)
- [Accounting Options Configuration on page 296](#)
- [Configuring the Interface Profile on page 309](#)
- [Configuring the Filter Profile on page 312](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level on page 295](#)

Configuring the Interface Profile

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

An interface profile specifies the information collected and written to a log file. You can configure a profile to collect error and statistic information for input and output packets on a particular physical or logical interface.

To configure an interface profile, include the **interface-profile** statement at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
interface-profile profile-name {
  fields {
    field-name;
  }
  file filename;
  interval minutes;
}
```

By default, the Packet Forwarding Engine (PFE) periodically collects the statistics for all interfaces. To improve the performance, you can optionally disable the periodic refresh by including the **periodic-refresh disable** statement at the **[edit accounting-options]** hierarchy level.

Each accounting profile must have a unique **profile-name**. To apply a profile to a physical or logical interface, include the **accounting-profile** statement at either the **[edit interfaces interface-name]** or the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. You can also apply an accounting profile at the **[edit firewall family family-type filter filter-name]** hierarchy level. For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

To configure an interface profile, perform the tasks described in the following sections:

- [Configuring Fields on page 310](#)
- [Configuring the File Information on page 310](#)
- [Configuring Cleared Statistics to be Reported in the Flat File on page 310](#)
- [Configuring the Interval on page 311](#)
- [Example: Configuring the Interface Profile on page 311](#)

Configuring Fields

An interface profile must specify what statistics are collected. To configure which statistics should be collected for an interface, include the **fields** statement at the **[edit accounting-options interface-profile profile-name]** hierarchy level:

```
[edit accounting-options interface-profile profile-name]
fields {
    field-name;
}
```

Configuring the File Information

Each accounting profile logs its statistics to a file in the **/var/log** directory.

To configure which file to use, include the **file** statement at the **[edit accounting-options interface-profile profile-name]** hierarchy level:

```
[edit accounting-options interface-profile profile-name]
file filename;
```

You must specify a **file** statement for the interface profile that has already been configured at the **[edit accounting-options]** hierarchy level.

Configuring Cleared Statistics to be Reported in the Flat File

When you issue the **clear interfaces statistics** command for a logical interface configured to collect accounting statistics, all accounting statistics received on that interface from the Packet Forwarding Engine are cleared. The current values when the command is issued become the new baseline and the statistics counters are reset to zero. The new values, starting from zero, are displayed in the CLI. However, they are not reported that way in the accounting flat file associated with the interface. Instead, the values as reported in the file continue to increment as if the command had not been issued.

You can change this result by including the **allow-clear** statement in the interface profile. In this case, when you issue the **clear interfaces statistics** command, the statistics are reset to zero and reported to the flat file.

To configure reporting of cleared accounting statistics to the flat file, specify reporting:

```
[edit accounting-options interface-profile profile-name]
allow-clear;
```

Configuring the Interval

Each interface with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the **interval** statement at the **[edit accounting-options interface-profile *profile-name*]** hierarchy level:

```
[edit accounting-options interface-profile profile-name]
interval minutes;
```



NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of interfaces might cause serious performance degradation.

The range for the **interval** statement is 1 through 2880 minutes. The default is 30 minutes.

Example: Configuring the Interface Profile

Configure the interface profile:

```
[edit]
accounting-options {
  file if_stats {
    size 40 files 5;
  }
  interface-profile if_profile1 {
    file if_stats;
    interval 30;
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-multicast;
      output-multicast;
    }
  }
  interface-profile if_profile2 {
    file if_stats;
    interval 30;
    fields {
      input-bytes;
      output-bytes;
      input-packets;
      output-packets;
      input-multicast;
      output-multicast;
    }
  }
}
```

```
interfaces {
  xe-1/0/0 {
    accounting-profile if_profile1;
    unit 0 {
      accounting-profile if_profile2;
      ...
    }
  }
}
```

The two interface profiles, if-profile1 and if-profile2, write data to the same file, if-stats. The if-stats file might look like the following:

```
#FILE CREATED 976823478 2000-12-14-19:51:18
#hostname host
#profile-layout
if_profile2,epoch-timestamp,interface-name,snmp-index,input-bytes,output-bytes,
input-packets,output-packets,input-multicast,output-multicast
#profile-layout
if_profile1,epoch-timestamp,interface-name,snmp-index,input-bytes,output-bytes,
input-packets
if_profile2,976823538,xe-1/0/0.0,8,134696815,3681534,501088,40723,0,0
if_profile1,976823538,xe-1/0/0,7,134696815,3681534,501088
...
#FILE CLOSED 976824378 2000-12-14-20:06:18
```

Related Documentation

- [Accounting Options Overview on page 291](#)
- [Understanding Device Management Functions in Junos OS on page 3](#)
- [Accounting Options Configuration on page 296](#)
- [Configuring Accounting-Data Log Files on page 304](#)
- [Configuring the Filter Profile on page 312](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level on page 295](#)

Configuring the Filter Profile

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

A filter profile specifies error and statistics information collected and written to a file. A filter profile must specify counter names for which statistics are collected.

To configure a filter profile, include the **filter-profile** statement at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
filter-profile profile-name {
  counters {
    counter-name;
  }
  file filename;
  interval minutes;
```

```
}
```

To apply the filter profile, include the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** hierarchy level.

To configure a filter profile, perform the tasks described in the following sections:

- [Configuring the Counters on page 313](#)
- [Configuring the File Information on page 313](#)
- [Configuring the Interval on page 314](#)

Configuring the Counters

Statistics are collected for all counters specified in the filter profile. To configure the counters, include the **counters** statement at the **[edit accounting-options filter-profile *profile-name*]** hierarchy level:

```
[edit accounting-options filter-profile profile-name]  
counters {  
}
```

Configuring the File Information

Each accounting profile logs its statistics to a file in the **/var/log** directory.

To configure which file to use, include the **file** statement at the **[edit accounting-options filter-profile *profile-name*]** hierarchy level:

```
[edit accounting-options filter-profile profile-name]  
file filename;
```

You must specify a filename for the filter profile that has already been configured at the **[edit accounting-options]** hierarchy level.



NOTE: The limit on the total number of characters per line in a log file equals 1023. If this limit is exceeded, the output written to the log file is incomplete. Ensure that you limit the number of counters or requested data so that this character limit is not exceeded.



NOTE: If the configured file size or transfer interval is exceeded, Junos OS closes the file and starts a new one. By default, the transfer interval value is 30 minutes. If the transfer interval is not configured, Junos OS closes the file and starts a new one when the file size exceeds its configured value or the default transfer interval value exceeds 30 minutes. To avoid transferring files every 30 minutes, specify a different value for the transfer interval.

Configuring the Interval

Each filter with an accounting profile enabled has statistics collected once per interval time specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the `interval` statement at the `[edit accounting-options filter-profile profile-name]` hierarchy level:

```
[edit accounting-options filter-profile profile-name]  
interval;
```



NOTE: The minimum interval allowed is 1 minute. Configuring a low interval in an accounting profile for a large number of filters might cause serious performance degradation.

The range for the `interval` statement is 1 through 2880 minutes. The default is 30 minutes.

Related Documentation

- [Accounting Options Overview on page 291](#)
- [Understanding Device Management Functions in Junos OS on page 3](#)
- [Accounting Options Configuration on page 296](#)
- [Configuring Accounting-Data Log Files on page 304](#)

Example: Configuring a Filter Profile

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Configure a filter profile:

```
[edit]  
accounting-options {  
  file fw_accounting {  
    size 500k files 4;  
  }  
  filter-profile fw_profile1 {  
    file fw_accounting;  
    interval 60;  
    counters {  
      counter1;  
      counter2;  
      counter3;  
    }  
  }  
}  
firewall {  
  filter myfilter {  
    accounting-profile fw_profile1;  
    ...  
    term accept-all {  
      then {  
        count counter1;  
        accept;  
      }  
    }  
  }  
}
```

```

    }
  }
}

```

The filter profile, **fw-profile1**, writes data to the file **fw_accounting**. The file might look like the following:

```

#FILE CREATED 976825278 2000-12-14-20:21:18
#hostname host
#profile-layout
fw_profile1,epoch-timestamp,filter-name,counter-name,packet-count,byte-count
fw_profile1,976826058,myfilter,counter1,163,10764
...
#FILE CLOSED 976826178 2000-12-14-20:36:18

```

Related Documentation

- [Configuring the Filter Profile on page 312](#)
- [Example: Configuring Interface-Specific Firewall Counters and Filter Profiles on page 315](#)

Example: Configuring Interface-Specific Firewall Counters and Filter Profiles

Supported Platforms [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

To collect and log count statistics collected by firewall filters on a per-interface basis, you must configure a filter profile and include the interface-specific statement at the **[edit firewall filter *filter-name*]** hierarchy level.

Configure the firewall filter accounting profile:

```

[edit accounting-options]
file cust1_accounting {
  size 500k;
}
filter-profile cust1_profile {
  file cust1_accounting;
  interval 1;
  counters {
    r1;
  }
}

```

Configure the interface-specific firewall counter:

```

[edit firewall]
filter f3 {
  accounting-profile cust1_profile;
  interface-specific;
  term f3-term {
    then {
      count r1;
      accept;
    }
  }
}

```

Apply the firewall filter to an interface:

```
[edit interfaces]
xe-1/0/0 {
  unit 0 {
    family inet {
      filter {
        input f3;
        output f3;
      }
      address 20.20.20.30/24;
    }
  }
}
```

The following example shows the contents of the **cust1_accounting** file in the **/var/log** folder that might result from the preceding configuration:

```
#FILE CREATED 995495212 2001-07-18-22:26:52
#hostname host
#profile-layout cust1_profile,epoch-timestamp,interfaces,filter-name,
counter-name,packet-count,byte-count
cust1_profile,995495572,xe-1/0/0.0,f3-xe-1/0/0.0-i,r1-xe-1/0/0.0-i,5953,1008257
cust1_profile,995495602,xe-1/0/0.0,f3-xe-1/0/0.0-o,r1-xe-1/0/0.0-o,5929,1006481
...
```

If the **interface-specific** statement is not included in the configuration, the following output might result:

```
#FILE CREATED 995495212 2001-07-18-22:26:52
#hostname host
#profile-layout cust1_profile,epoch-timestamp,interfaces,filter-name,
counter-name,packet-count,byte-count
cust1_profile,995495572,xe-1/0/0.0,f3,r1,5953,1008257
cust1_profile,995495632,xe-1/0/0.0,f3,r1,5929,1006481
```

- Related Documentation**
- [Configuring the Filter Profile on page 312](#)
 - [Configuring the Interface Profile on page 309](#)

Configuring SCU or DCU

Supported Platforms **M Series, MX Series, PTX Series, SRX Series, T Series, vSRX**

To configure SCU or DCU, perform the following tasks described in this section:



NOTE: We recommend that you stop the network traffic on an interface before you modify the DCU or SCU configuration for that interface. Modifying the DCU or SCU configuration without stopping the traffic might corrupt the DCU or SCU statistics. Before you restart the traffic after modifying the configuration, enter the clear interfaces statistics command.

- [Creating Prefix Route Filters in a Policy Statement on page 317](#)
- [Applying the Policy to the Forwarding Table on page 317](#)
- [Enabling Accounting on Inbound and Outbound Interfaces on page 317](#)

Creating Prefix Route Filters in a Policy Statement

To define prefix router filters:

```
[edit policy-options]
policy-statement scu-1 {
  term term1;
  from {
    route-filter 192.0.2.0/24 or longer;
  }
  then source-class gold;
}
```

Applying the Policy to the Forwarding Table

To apply the policy to the forwarding table:

```
[edit]
routing-options {
  forwarding-table {
    export scu-1;
  }
}
```

Enabling Accounting on Inbound and Outbound Interfaces

To enable accounting on inbound and outbound interfaces:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet;
      accounting {
        destination-class-usage;
        source-class-usage {
          output;
        }
      }
    }
  }
}
[edit]
```

```
interfaces {
  xe-0/1/0 {
    unit 0 {
      family inet6 {
        accounting {
          source-class-usage {
            input;
          }
        }
      }
    }
  }
}
```

Optionally, you can include the input and output statements on a single interface as shown:

```
[edit]
interfaces {
  xe-0/1/2 {
    unit 0 {
      family inet6 {
        accounting {
          source-class-usage {
            input;
            output;
          }
        }
      }
    }
  }
}
```

For more information about configuring route filters and source classes in a routing policy, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide* and the *Junos OS Network Interfaces Library for Routing Devices*.

**Related
Documentation**

- [Understanding Source Class Usage and Destination Class Usage Options on page 292](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 318](#)
- [Configuring Class Usage Profiles on page 320](#)
- [Configuring the MIB Profile on page 322](#)
- [Configuring the Routing Engine Profile on page 325](#)

Configuring SCU on a Virtual Loopback Tunnel Interface

Supported Platforms M Series, MX Series, SRX Series, T Series, vSRX

To configure source class usage on the virtual loopback tunnel interface, perform the tasks described in the following sections:

- [Example: Configuring a Virtual Loopback Tunnel Interface on a Provider Edge Router Equipped with a Tunnel PIC on page 319](#)
- [Example: Mapping the VRF Instance Type to the Virtual Loopback Tunnel Interface on page 319](#)
- [Example: Sending Traffic Received from the Virtual Loopback Interface Out the Source Class Output Interface on page 320](#)

Example: Configuring a Virtual Loopback Tunnel Interface on a Provider Edge Router Equipped with a Tunnel PIC

Define a virtual loop interface on a provider edge router with a Tunnel PIC:

```
[edit interfaces]
vt-0/3/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          input;
        }
      }
    }
  }
}
```

Example: Mapping the VRF Instance Type to the Virtual Loopback Tunnel Interface

Map the VRF instance type to the virtual loopback tunnel interface:

```
[edit]
routing-instances {
  VPN-A {
    instance-type vrf;
    interface at-2/1/1.0;
    interface vt-0/3/0.0;
    route-distinguisher 10.255.14.225;
    vrf-import import-policy-name;
    vrf-export export-policy-name;
    protocols {
      bgp {
        group to-r4 {
          local-address 10.27.253.1;
          peer-as 400;
          neighbor 10.27.253.2;
        }
      }
    }
  }
}
```



NOTE: For SCU and DCU to work, do not include the `vrf-table-label` statement at the `[edit routing-instances instance-name]` hierarchy level.

Example: Sending Traffic Received from the Virtual Loopback Interface Out the Source Class Output Interface

Send traffic received from the virtual loopback tunnel interface out of the source class output interface:

```
[edit interfaces]
at-1/1/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          output;
        }
      }
    }
  }
}
```

For more information about configuring source class usage on the virtual loopback tunnel interface, see the *Junos OS Network Interfaces Library for Routing Devices*.

Related Documentation

- [Understanding Source Class Usage and Destination Class Usage Options on page 292](#)
- [Configuring SCU or DCU on page 316](#)
- [Configuring Class Usage Profiles on page 320](#)
- [Configuring the MIB Profile on page 322](#)
- [Configuring the Routing Engine Profile on page 325](#)

Configuring Class Usage Profiles

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

To collect class usage statistics, perform the tasks described in these sections:

- [Configuring a Class Usage Profile on page 320](#)
- [Configuring the File Information on page 321](#)
- [Configuring the Interval on page 321](#)
- [Creating a Class Usage Profile to Collect Source Class Usage Statistics on page 321](#)
- [Creating a Class Usage Profile to Collect Destination Class Usage Statistics on page 322](#)

Configuring a Class Usage Profile

You can configure the class usage profile to collect statistics for particular source and destination classes.

To configure the class usage profile to filter by source classes, include the **source-classes** statement at the **[edit accounting-options class-usage-profile *profile-name*]** hierarchy level:

```
[edit accounting-options class-usage-profile profile-name]  
source-classes {  
    source-class-name;  
}
```

To configure the class usage profile to filter by destination classes, include the **destination-classes** statement at the **[edit accounting-options class-usage-profile *profile-name*]** hierarchy level:

```
[edit accounting-options class-usage-profile profile-name]  
destination-classes {  
    destination-class-name;  
}
```

Configuring the File Information

Each accounting profile logs its statistics to a file in the **/var/log** directory.

To specify which file to use, include the **file** statement at the **[edit accounting-options class-usage-profile *profile-name*]** hierarchy level:

```
[edit accounting-options class-usage-profile profile-name]  
file filename;
```

You must specify a filename for the source class usage profile that has already been configured at the **[edit accounting-options]** hierarchy level. You can also specify a filename for the destination class usage profile configured at the **[edit accounting-options]** hierarchy level.

Configuring the Interval

Each interface with a class usage profile enabled has statistics collected once per interval specified for the accounting profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the **interval** statement at the **[edit accounting-options class-usage-profile *profile-name*]** hierarchy level:

```
[edit accounting-options class-usage-profile profile-name]  
interval;
```

Creating a Class Usage Profile to Collect Source Class Usage Statistics

To create a class usage profile to collect source class usage statistics:

```
[edit]  
accounting-options {  
    class-usage-profile scu-profile;  
    file usage-stats;  
    interval 15;  
    source-classes {  
        gold;  
        silver;  
        bronze;  
    }  
}
```

The class usage profile, **scu-profile1**, writes data to the file **usage_stats**. The file might look like the following:

```
#FILE CREATED 976825278 2000-12-14-20:21:18
#profile-layout, scu_profile,epoch-timestamp,interface-name,source-class,
packet-count,byte-count
scu_profile,980313078,xe-1/0/0.0,gold,82,6888
scu_profile,980313078,xe-1/0/0.0,silver,164,13776
scu_profile,980313078,xe-1/0/0.0,bronze,0,0
scu_profile,980313678,xe-1/0/0.0,gold,82,6888
scu_profile,980313678,xe-1/0/0.0,silver,246,20664
scu_profile,980313678,xe-1/0/0.0,bronze,0,0
```

Creating a Class Usage Profile to Collect Destination Class Usage Statistics

To create a class usage profile to collect destination class usage statistics:

```
[edit]
accounting-options {
  class-usage-profile dcu-profile1;
  file usage-stats
  interval 15;
  destination-classes {
    gold;
    silver;
    bronze;
  }
}
```

The class usage profile, **dcu-profile1**, writes data to the file **usage_stats**. The file might look like the following:

```
#FILE CREATED 976825278 2000-12-14-20:21:18
#profile-layout, dcu_profile,epoch-timestamp,interface-name,destination-class,
packet-count,byte-count
dcu_profile,980313078,xe-1/0/0.0,gold,82,6888
dcu_profile,980313078,xe-1/0/0.0,silver,164,13776
dcu_profile,980313078,xe-1/0/0.0,bronze,0,0
dcu_profile,980313678,xe-1/0/0.0,gold,82,6888
dcu_profile,980313678,xe-1/0/0.0,silver,246,20664
dcu_profile,980313678,xe-1/0/0.0,bronze,0,0
...
#FILE CLOSED 976826178 2000-12-14-20:36:18
```

Related Documentation

- [Understanding Source Class Usage and Destination Class Usage Options on page 292](#)
- [Configuring SCU or DCU on page 316](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 318](#)
- [Configuring the Routing Engine Profile on page 325](#)

Configuring the MIB Profile

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

The MIB profile collects MIB statistics and logs them to a file. The MIB profile specifies the SNMP operation and MIB object names for which statistics are collected.

To configure a MIB profile, include the **mib-profile** statement at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
mib-profile profile-name {
  file filename;
  interval minutes;
  object-names {
    mib-object-name;
  }
  operation operation-name;
}
```

To configure a MIB profile, perform the tasks described in the following sections:

- [Configuring the File Information on page 323](#)
- [Configuring the Interval on page 323](#)
- [Configuring the MIB Operation on page 323](#)
- [Configuring MIB Object Names on page 324](#)
- [Example: Configuring a MIB Profile on page 324](#)

Configuring the File Information

Each accounting profile logs its statistics to a file in the **/var/log** directory.

To configure which file to use, include the **file** statement at the **[edit accounting-options mib-profile profile-name]** hierarchy level:

```
[edit accounting-options mib-profile profile-name]
file filename;
```

You must specify a **filename** for the MIB profile that has already been configured at the **[edit accounting-options]** hierarchy level.

Configuring the Interval

A MIB profile has statistics collected once per interval time specified for the profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the **interval** statement at the **[edit accounting-options mib-profile profile-name]** hierarchy level:

```
[edit accounting-options mib-profile profile-name]
interval;
```

The range for the **interval** statement is 1 through 2880 minutes. The default is 30 minutes.

Configuring the MIB Operation

A MIB profile must specify the operation that is used to collect MIB statistics. To configure which operation is used to collect MIB statistics, include the **operation** statement at the **[edit accounting-options mib-profile profile-name]** hierarchy level:

```
[edit accounting-options mib-profile profile-name]
```

operation *operation-name*;

You can configure a **get**, **get-next**, or **walk** operation. The default operation is **walk**.

Configuring MIB Object Names

A MIB profile must specify the MIB objects for which statistics are to be collected. To configure the MIB objects for which statistics are collected, include the **objects-names** statement at the **[edit accounting-options mib-profile *profile-name*]** hierarchy level:

```
[edit accounting-options mib-profile profile-name]
object-names {
  mib-object-name;
}
```

You can include multiple MIB object names in the configuration.



NOTE: In Junos OS Release 15.1X49-D10 and later, do not configure MIB objects related to interface octets or packets for a MIB profile, because it can cause the SNMP walk or a CLI show command to time out.

Example: Configuring a MIB Profile

Configure a MIB profile:

```
[edit accounting-options]
mib-profile mstatistics {
  file stats;
  interval 60;
  operation walk;
  objects-names {
    ipCidrRouteStatus;
  }
}
```

Release History Table

Release	Description
15.1X49-D10	In Junos OS Release 15.1X49-D10 and later, do not configure MIB objects related to interface octets or packets for a MIB profile, because it can cause the SNMP walk or a CLI show command to time out.

Related Documentation

- [Understanding Source Class Usage and Destination Class Usage Options on page 292](#)
- [Configuring SCU or DCU on page 316](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 318](#)
- [Configuring Class Usage Profiles on page 320](#)
- [Configuring the Routing Engine Profile on page 325](#)

Configuring the Routing Engine Profile

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

The Routing Engine profile collects Routing Engine statistics and logs them to a file. The Routing Engine profile specifies the fields for which statistics are collected.

To configure a Routing Engine profile, include the **routing-engine-profile** statement at the **[edit accounting-options]** hierarchy level:

```
[edit accounting-options]
routing-engine-profile profile-name {
  fields {
    field-name;
  }
  file filename;
  interval minutes;
}
```

To configure a Routing Engine profile, perform the tasks described in the following sections:

- [Configuring Fields on page 325](#)
- [Configuring the File Information on page 325](#)
- [Configuring the Interval on page 326](#)
- [Example: Configuring a Routing Engine Profile on page 326](#)

Configuring Fields

A Routing Engine profile must specify what statistics are collected. To configure which statistics should be collected for the Routing Engine, include the **fields** statement at the **[edit accounting-options routing-engine-profile profile-name]** hierarchy level:

```
[edit accounting-options routing-engine-profile profile-name]
fields {
  field-name;
}
```

Configuring the File Information

Each accounting profile logs its statistics to a file in the **/var/log** directory.

To configure which file to use, include the **file** statement at the **[edit accounting-options routing-engine-profile profile-name]** hierarchy level:

```
[edit accounting-options routing-engine-profile profile-name]
file filename;
```

You must specify a **filename** for the Routing Engine profile that has already been configured at the **[edit accounting-options]** hierarchy level.

Configuring the Interval

A Routing Engine profile has statistics collected once per interval time specified for the profile. Statistics collection time is scheduled evenly over the configured interval. To configure the interval, include the **interval** statement at the **[edit accounting-options routing-engine-profile *profile-name*]** hierarchy level:

```
[edit accounting-options routing-engine-profile profile-name]  
interval;
```

The range for **interval** is 1 through 2880 minutes. The default is 30 minutes.

Example: Configuring a Routing Engine Profile

Configure a Routing Engine profile:

```
[edit accounting-options]  
file my-file {  
    size 300k;  
}  
routing-engine-profile profile-1 {  
    file my-file;  
    fields {  
        host-name;  
        date;  
        time-of-day;  
        uptime;  
        cpu-load-1;  
        cpu-load-5;  
        cpu-load-15;  
    }  
}
```

Related Documentation

- [Understanding Source Class Usage and Destination Class Usage Options on page 292](#)
- [Configuring SCU or DCU on page 316](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 318](#)
- [Configuring Class Usage Profiles on page 320](#)
- [Configuring the MIB Profile on page 322](#)

PART 6

Configuring Monitoring Options

- [Configuring Interface Alarms on page 329](#)
- [Using RPM to Measure Network Performance on page 341](#)
- [Configuring IP Monitoring on page 371](#)

Configuring Interface Alarms

- [Alarm Overview on page 329](#)
- [Example: Configuring Interface Alarms on page 335](#)
- [Monitoring Active Alarms on a Device on page 338](#)
- [Monitoring Alarms on page 339](#)

Alarm Overview

Supported Platforms [SRX Series](#)

Alarms alert you to conditions on a network interface, on the device chassis, or in the system software that might prevent the device from operating normally. You can set the conditions that trigger alarms on an interface. Chassis and system alarm conditions are preset.

An active alarm lights the **ALARM** LED on the front panel of the device. You can monitor active alarms from the J-Web user interface or the CLI. When an alarm condition triggers an alarm, the device lights the yellow (amber) **ALARM** LED on the front panel. When the condition is corrected, the light turns off.

This section contains the following topics:

- [Alarm Types on page 329](#)
- [Alarm Severity on page 330](#)
- [Alarm Conditions on page 330](#)

Alarm Types

The device supports three types of alarms:

- Interface alarms indicate a problem in the state of the physical links on fixed or installed Physical Interface Modules (PIMs). To enable interface alarms, you must configure them.
- Chassis alarms indicate a failure on the device or one of its components. Chassis alarms are preset and cannot be modified.

- System alarms indicate a missing rescue configuration or software license, where valid. System alarms are preset and cannot be modified, although you can configure them to appear automatically in the J-Web user interface or CLI.

Starting with Junos OS Release 15.1X49-D60 and Junos OS Release 17.3R1, a new system alarm is introduced to indicate that the PICs (I/O card or SPC) have failed to come online during system start time.



NOTE: Run the following commands when the CLI prompt indicates that an alarm has been raised:

- `show system alarms`
- `show chassis alarms`
- `show chassis fpc pic-status`

For more information about the CLI commands, see [show system alarms](#), [show chassis alarms](#), and [show chassis fpc \(View\)](#).

Alarm Severity

Alarms have two severity levels:

- Major (red)—Indicates a critical situation on the device that has resulted from one of the following conditions. A red alarm condition requires immediate action.
 - One or more hardware components have failed.
 - One or more hardware components have exceeded temperature thresholds.
 - An alarm condition configured on an interface has triggered a critical warning.
- Minor (yellow)—Indicates a noncritical condition on the device that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

A missing rescue configuration or software license generates a yellow system alarm.

Alarm Conditions

To enable alarms on a device interface, you must select an alarm condition and an alarm severity. In contrast, alarm conditions and severity are preconfigured for chassis alarms and system alarms.



NOTE: For information about chassis alarms for your device, see the Hardware Guide for your device.

This section contains the following topics:

- [Interface Alarm Conditions on page 331](#)
- [System Alarm Conditions on page 334](#)

Interface Alarm Conditions

[Table 38 on page 331](#) lists the interface conditions, sorted by interface type, that you can configure for an alarm. You can configure each alarm condition to trigger either a major (red) alarm or minor a (yellow) alarm. The corresponding configuration option is included.

For the services stateful firewall filters (NAT, IDP, and IPsec), which operate on an internal adaptive services module within a device, you can configure alarm conditions on the integrated services and services interfaces.

Table 38: Interface Alarm Conditions

Interface	Alarm Condition	Description	Configuration Option
DS1 (T1)	Alarm indication signal (AIS)	The normal T1 traffic signal contained a defect condition and has been replaced by the AIS. A transmission interruption occurred at the remote endpoint or upstream of the remote endpoint. This all-ones signal is transmitted to prevent consequential downstream failures or alarms.	ais
	Yellow alarm	The remote endpoint is in yellow alarm failure. This condition is also known as a far-end alarm failure.	ylw
Ethernet	Link is down	The physical link is unavailable.	link-down
Integrated services	Hardware or software failure	On the adaptive services module, either the hardware associated with the module or the software that drives the module has failed.	failure

Table 38: Interface Alarm Conditions (*continued*)

Interface	Alarm Condition	Description	Configuration Option
Serial	Clear-to-send (CTS) signal absent	The remote endpoint of the serial link is not transmitting a CTS signal. The CTS signal must be present before data can be transmitted across a serial link.	cts-absent
	Data carrier detect (DCD) signal absent	The remote endpoint of the serial link is not transmitting a DCD signal. Because the DCD signal transmits the state of the device, no signal probably indicates that the remote endpoint of the serial link is unavailable.	dcd-absent
	Data set ready (DSR) signal absent	The remote endpoint of the serial link is not transmitting a DSR signal. The DSR signal indicates that the remote endpoint is ready to receive and transmit data across the serial link.	dsr-absent
	Loss of receive clock	The clock signal from the remote endpoint is not present. Serial connections require clock signals to be transmitted from one endpoint and received by the other endpoint of the link.	loss-of-rx-clock
	Loss of transmit clock	The local clock signal is not present. Serial connections require clock signals to be transmitted from one endpoint and received by the other endpoint of the link.	loss-of-tx-clock
Services	Services module hardware down	A hardware problem has occurred on the device's services module. This error typically means that one or more of the CPUs on the module has failed.	hw-down
	Services link down	The link between the device and its services module is unavailable.	linkdown
	Services module held in reset	The device's services module is stuck in reset mode. If the services module fails to start up five or more times in a row, the services module is held in reset mode. Startup fails when the amount of time from CPU release to CPU halt is less than 300 seconds.	pic-hold-reset
	Services module reset	The device's services module is resetting. The module resets after it crashes or is reset from the CLI, or when it takes longer than 60 seconds to start up.	pic-reset
	Services module software down	A software problem has occurred on the device's services module.	sw-down

Table 38: Interface Alarm Conditions (*continued*)

Interface	Alarm Condition	Description	Configuration Option
E3	Alarm indication signal (AIS)	The normal E3 traffic signal contained a defect condition and has been replaced by the AIS. A transmission interruption occurred at the remote endpoint or upstream of the remote endpoint. This all-ones signal is transmitted to prevent consequential downstream failures or alarms.	ais
	Loss of signal (LOS)	No remote E3 signal is being received at the E3 interface.	los
	Out of frame (OOF)	An OOF condition has existed for 10 seconds. This alarm applies only to E3 interfaces configured in frame mode. The OOF failure is cleared when no OOF or LOS defects have occurred for 20 seconds.	oof
	Remote defect indication	An AIS, LOS, or OOF condition exists. This alarm applies only to E3 interfaces configured in frame mode.	rdi

Table 38: Interface Alarm Conditions (*continued*)

Interface	Alarm Condition	Description	Configuration Option
T3 (DS3)	Alarm indication signal	The normal T3 traffic signal contained a defect condition and has been replaced by the AIS. A transmission interruption occurred at the remote endpoint or upstream of the remote endpoint. This all-ones signal is transmitted to prevent consequential downstream failures or alarms.	ais
	Excessive number of zeros	The bit stream received from the upstream host has more consecutive zeros than are allowed in a T3 frame.	exz
	Far-end receive failure (FERF)	The remote endpoint of the connection has failed. A FERF differs from a yellow alarm, because the failure can be any failure, not just an OOF or LOS failure.	ferf
	Idle alarm	The Idle signal is being received from the remote endpoint.	idle
	Line code violation	Either the line encoding along the T3 link is corrupted or a mismatch between the encoding at the local and remote endpoints of a T3 connection occurred.	lcv
	Loss of frame (LOF)	An OOF or loss-of-signal LOS condition has existed for 10 seconds. The LOF failure is cleared when no OOF or LOS defects have occurred for 20 seconds. A LOF failure is also called a red failure.	lof
	Loss of signal (LOS)	No remote T3 signal is being received at the T3 interface.	los
	Phase-locked loop out of lock	The clocking signals for the local and remote endpoints no longer operate in lock-step.	pll
	Yellow alarm	The remote endpoint is in yellow alarm failure. This condition is also known as a far-end alarm failure.	ylw

System Alarm Conditions

Table 39 on page 334 lists the two preset system alarms, the condition that triggers each alarm, and the action you take to correct the condition.

Table 39: System Alarm Conditions and Corrective Actions

Alarm Type	Alarm Condition	Corrective Action
Configuration	The rescue configuration is not set.	Set the rescue configuration.

Table 39: System Alarm Conditions and Corrective Actions (*continued*)

Alarm Type	Alarm Condition	Corrective Action
License	<p>You have configured at least one software feature that requires a feature license, but no valid license for the feature is currently installed.</p> <p>NOTE: This alarm indicates that you are in violation of the software license agreement. You must install a valid license key to be in compliance with all agreements.</p>	Install a valid license key.

Release History Table

Release	Description
15.1X49-D60	Starting with Junos OS Release 15.1X49-D60 and Junos OS Release 17.3R1, a new system alarm is introduced to indicate that the PICs (I/O card or SPC) have failed to come online during system start time.

Related Documentation

- [Example: Configuring Interface Alarms on page 335](#)
- [Monitoring Active Alarms on a Device on page 338](#)
- [Monitoring Alarms on page 339](#)
- [System Log Messages](#)

Example: Configuring Interface Alarms**Supported Platforms** [SRX Series](#)

This example shows how to configure interface alarms.

- [Requirements on page 335](#)
- [Overview on page 336](#)
- [Configuration on page 336](#)
- [Verification on page 338](#)

Requirements

Before you begin:

- Establish basic connectivity.
- Configure network interfaces. See [Interfaces Feature Guide for Security Devices](#).
- Select the network interface on which to apply an alarm and the condition you want to trigger the alarm. See [“Alarm Overview” on page 329](#).

Overview

In this example, you enable interface alarms by explicitly setting alarm conditions. You configure the system to generate a red interface alarm when a yellow alarm is detected on a DS1 link. You configure the system to generate a red interface alarm when a link-down failure is detected on an Ethernet link.

For a serial link, you set `cts-absent` and `dcd-absent` to yellow to signify either the CST or the DCD signal is not detected. You set `loss-of-rx-clock` and `loss-of-tx-clock` to red alarm to signify either the receiver clock signal or the transmission clock signal is not detected.

For a T3 link, you set the interface alarm to red when the remote endpoint is experiencing a failure. You set `exz` to yellow alarm when the upstream bit has more consecutive zeros than are permitted in a T3 interface. You then set a red alarm when there is loss-of-signal on the interface.

Finally, you configure the system to display active system alarms whenever a user with the login class `admin` logs in to the device.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis alarm ds1 ylw red
set chassis alarm ethernet link-down red
set chassis alarm serial cts-absent yellow dcd-absent yellow
set chassis alarm serial loss-of-rx-clock red loss-of-tx-clock red
set chassis alarm t3 ylw red exz yellow los red
set system login class admin login-alarms
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure interface alarms:

1. Configure an alarm.

```
[edit]
user@host# edit chassis alarm
```
2. Specify the interface alarms on a DS1 and an Ethernet link.

```
[edit chassis alarm]
user@host# set ds1 ylw red
user@host# set ethernet link-down red
```
3. Specify the interface alarms on a serial link.

```
[edit chassis alarm]
user@host# set serial cts-absent yellow
user@host# set serial dcd-absent yellow
user@host# set serial loss-of-rx-clock red
user@host# set serial loss-of-tx-clock red
```

4. Specify the interface alarms on a T3 link.

```
[edit chassis alarm]
user@host# set t3 ylw red
user@host# set t3 exz yellow
user@host# set t3 los red
```

5. Configure the system to display active system alarms.

```
[edit]
user@host# edit system login
user@host# set class admin login-alarms
```

Results From configuration mode, confirm your configuration by entering the **show chassis alarms** and **show system login** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show chassis alarms
t3 {
  exz yellow;
  los red;
  ylw red;
}
ds1 {
  ylw red;
}
ethernet {
  link-down red;
}
serial {
  loss-of-rx-clock red;
  loss-of-tx-clock red;
  dcd-absent yellow;
  cts-absent yellow;
}
[edit]
user@host# show system login
show system login
show system login
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Alarm Configurations

Purpose	Confirm that the configuration is working properly. Verify that the alarms are configured.
Action	From configuration mode, enter the show chassis alarms command. Verify that the output shows the intended configuration of the alarms.
Related Documentation	<ul style="list-style-type: none">• Alarm Overview on page 329• Monitoring Active Alarms on a Device on page 338• Monitoring Alarms on page 339

Monitoring Active Alarms on a Device

Supported Platforms [SRX Series, vSRX](#)

Purpose	Use to monitor and filter alarms on a Juniper Networks device.
Action	<p>Select Monitor>Events and Alarms>View Alarms in the J-Web user interface. The J-Web View Alarms page displays the following information about preset system and chassis alarms:</p> <ul style="list-style-type: none">• Type—Type of alarm: System, Chassis, or All.• Severity—Severity class of the alarm: Minor or Major.• Description—Description of the alarm.• Time—Time that the alarm was registered. <p>To filter which alarms appear, use the following options:</p> <ul style="list-style-type: none">• Alarm Type—Specifies which type of alarm to monitor: System, Chassis, or All. System alarms include FRU detection alarms (power supplies removed, for instance). Chassis alarms indicate environmental alarms such as temperature.• Severity—Specifies the alarm severity that you want to monitor: Major, Minor, or All. A major (red) alarm condition requires immediate action. A minor (yellow) condition requires monitoring and maintenance.• Description—Specifies the alarms you want to monitor. Enter a brief synopsis of the alarms that you want to monitor.• Date From—Specifies the beginning of the date range that you want to monitor. Set the date using the calendar pick tool.

- **To**—Specifies the end of the date range that you want to monitor. Set the date using the calendar pick tool.
- **Go**—Executes the options that you specified.
- **Reset**—Clears the options that you specified.

Alternatively, you can enter the following **show** commands in the CLI editor:

- **show chassis alarms**
- **show system alarms**

Related Documentation

- [Alarm Overview on page 329](#)
- [Example: Configuring Interface Alarms on page 335](#)
- [Monitoring Alarms on page 339](#)

Monitoring Alarms

Supported Platforms [SRX Series, vSRX](#)

Purpose Use the monitoring functionality to view the alarms page.

Action To monitor alarms, select one of the following in the J-Web user interface:

- If you are using SRX5400, SRX5600, or SRX5800 platforms, select **Monitor>Events and Alarms>View Alarms**.
- Select **Monitor>Alarms>View Alarms**.

Meaning [Table 40 on page 339](#) summarizes key output fields in the alarms page.

Table 40: Alarms Monitoring Page

Field	Value	Additional Information
Alarm Filter		
Alarm Type	Specifies the type of alarm to monitor: <ul style="list-style-type: none"> • System— System alarms include FRU detection alarms (power supplies removed, for instance). • Chassis— Chassis alarms indicate environmental alarms such as temperature. • All— Indicates to display all the types of alarms. 	—

Table 40: Alarms Monitoring Page (*continued*)

Field	Value	Additional Information
Severity	<p>Specifies the alarm severity that you want to monitor</p> <ul style="list-style-type: none"> • Major— A major (red) alarm condition requires immediate action. • Minor— A minor (yellow) condition requires monitoring and maintenance. • All— Indicates to display all the severities. 	—
Description	Enter a brief synopsis of the alarms you want to monitor.	—
Date From	Specifies the beginning of the date range that you want to monitor. Set the date using the calendar pick tool.	—
To	Specifies the end of the date range that you want to monitor. Set the date using the calendar pick tool.	—
Go	Executes the options that you specified.	—
Reset	Clears the options that you specified.	—
Alarm Details	<p>Displays the following information about each alarm:</p> <ul style="list-style-type: none"> • Type— Type of alarm: System, Chassis, or All. • Severity— Severity class of the alarm: Minor or Major. • Description— Description of the alarm. • Time— Time that the alarm was registered. 	—

- Related Documentation**
- [Monitoring Active Alarms on a Device on page 338](#)
 - [Monitoring Events on page 480](#)
 - [Monitoring Security Events by Policy on page 462](#)

CHAPTER 20

Using RPM to Measure Network Performance

- [RPM Overview on page 341](#)
- [IPv6 RPM Probes on page 345](#)
- [Two-Way Active Measurement Protocol \(TWAMP\) Overview on page 346](#)
- [Guidelines for Configuring RPM Probes for IPv6 on page 348](#)
- [RPM Support for VPN Routing and Forwarding on page 349](#)
- [Example: Configuring Basic RPM Probes on page 350](#)
- [Example: Configuring RPM Using TCP and UDP Probes on page 354](#)
- [Example: Configuring RPM Probes for BGP Monitoring on page 357](#)
- [Directing RPM Probes to Select BGP Devices on page 360](#)
- [Configuring IPv6 RPM Probes on page 360](#)
- [Tuning RPM Probes on page 361](#)
- [RPM Configuration Options on page 362](#)
- [Monitoring RPM Probes on page 366](#)

RPM Overview

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

The real-time performance monitoring (RPM) feature allows network operators and their customers to accurately measure the performance between two network endpoints. With the RPM tool, you configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter.

RPM allows you to perform service-level monitoring. When RPM is configured on a device, the device calculates network performance based on packet response time, jitter, and packet loss. These values are gathered by Hypertext Transfer Protocol (HTTP) GET requests, Internet Control Message Protocol (ICMP) requests, and TCP and UDP requests, depending on the configuration.

This section contains the following topics:

- [RPM Probes on page 342](#)
- [RPM Tests on page 343](#)
- [Probe and Test Intervals on page 343](#)
- [Jitter Measurement with Hardware Timestamping on page 343](#)
- [RPM Statistics on page 344](#)
- [RPM Thresholds and Traps on page 345](#)
- [RPM for BGP Monitoring on page 345](#)

RPM Probes

You gather RPM statistics by sending out probes to a specified probe target, identified by an IP address or URL. When the target receives the probe, it generates responses, which are received by the device. By analyzing the transit times to and from the remote server, the device can determine network performance statistics.

The device sends out the following probe types:

- HTTP GET request at a target URL
- HTTP GET request for metadata at a target URL
- ICMP echo request to a target address (the default)
- ICMP timestamp request to a target address
- UDP ping packets to a target device
- UDP timestamp requests to a target address
- TCP ping packets to a target device

UDP and TCP probe types require that the remote server be configured as an RPM receiver so that it generates responses to the probes.

The RPM probe results are also available in the form of MIB objects through the SNMP protocol.



NOTE: On SRX300, SRX320, SRX340, SRX1500 devices and vSRX instances, when you configure basic RPM probes, the following combination of the configuration parameters is not supported:

Source address and destination port and next-hop.

Configuring RPM probe with these parameters prevents sending out RPM probes to a specified probe target. We recommend you to configure either the source address or destination port and next-hop to configure RPM probe.

RPM Tests

Each probed target is monitored over the course of a test. A test represents a collection of probes, sent out at regular intervals, as defined in the configuration. Statistics are then returned for each test. Because a test is a collection of probes that have been monitored over some amount of time, test statistics such as standard deviation and jitter can be calculated and included with the average probe statistics.

Probe and Test Intervals

Within a test, RPM probes are sent at regular intervals, configured in seconds. When the total number of probes has been sent and the corresponding responses received, the test is complete. You can manually set the probe interval for each test to control how the RPM test is conducted.

After all the probes for a particular test have been sent, the test begins again. The time between tests is the test interval. You can manually set the test interval to tune RPM performance.



NOTE: On SRX340 devices, the RPM server operation with icmp is not supported. The RPM server works fine with TCP and UDP.

Jitter Measurement with Hardware Timestamping

Jitter is the difference in relative transit time between two consecutive probes.

You can timestamp the following RPM probes to improve the measurement of latency or jitter:

- ICMP ping
- ICMP ping timestamp
- UDP ping
- UDP ping timestamp



NOTE: The device supports hardware timestamping of UDP ping and UDP ping timestamp RPM probes only if the destination port is UDP-ECHO (port 7).

Timestamping takes place during the forwarding process of the device originating the probe (the RPM client), but not on the remote device that is the target of the probe (the RPM server).

The supported encapsulations on a device for timestamping are Ethernet including VLAN, synchronous PPP, and Frame Relay. The only logical interface supported is an `lt` services interface.

RPM probe generation with hardware timestamp can be retrieved through the SNMP protocol.

RPM Statistics

At the end of each test, the device collects the statistics for packet round-trip time, packet inbound and outbound times (for ICMP timestamp probes only), and probe loss as shown in [Table 41 on page 344](#).

Table 41: RPM Statistics

RPM Statistics	Description
Round-Trip Times	
Minimum round-trip time	Shortest round-trip time from the Juniper Networks device to the remote server, as measured over the course of the test
Maximum round-trip time	Longest round-trip time from the Juniper Networks device to the remote server, as measured over the course of the test
Average round-trip time	Average round-trip time from the Juniper Networks device to the remote server, as measured over the course of the test
Standard deviation round-trip time	Standard deviation of the round-trip times from the Juniper Networks device to the remote server, as measured over the course of the test
Jitter	Difference between the maximum and minimum round-trip times, as measured over the course of the test
Inbound and Outbound Times (ICMP Timestamp Probes Only)	
Minimum egress time	Shortest one-way time from the Juniper Networks device to the remote server, as measured over the course of the test
Maximum ingress time	Shortest one-way time from the remote server to the Juniper Networks device, as measured over the course of the test
Average egress time	Average one-way time from the Juniper Networks device to the remote server, as measured over the course of the test
Average ingress time	Average one-way time from the remote server to the Juniper Networks device, as measured over the course of the test
Standard deviation egress time	Standard deviation of the one-way times from the Juniper Networks device to the remote server, as measured over the course of the test
Standard deviation ingress time	Standard deviation of the one-way times from the remote server to the Juniper Networks device, as measured over the course of the test
Egress jitter	Difference between the maximum and minimum outbound times, as measured over the course of the test
Ingress jitter	Difference between the maximum and minimum inbound times, as measured over the course of the test

Table 41: RPM Statistics (*continued*)

RPM Statistics	Description
Probe Counts	
Probes sent	Total number of probes sent over the course of the test
Probe responses received	Total number of probe responses received over the course of the test
Loss percentage	Percentage of probes sent for which a response was not received

RPM Thresholds and Traps

You can configure RPM threshold values for the round-trip times, ingress (inbound) times, and egress (outbound) times that are measured for each probe, as well as for the standard deviation and jitter values that are measured for each test. Additionally, you can configure threshold values for the number of successive lost probes within a test and the total number of lost probes within a test.

If the result of a probe or test exceeds any threshold, the device generates a system log message and sends any Simple Network Management Protocol (SNMP) notifications (traps) that you have configured.

RPM for BGP Monitoring

When managing peering networks that are connected using Border Gateway Protocol (BGP), you might need to find out if a path exists between the Juniper Networks device and its configured BGP neighbors. You can ping each BGP neighbor manually to determine the connection status, but this method is not practical when the device has a large number of BGP neighbors configured.

In the device, you can configure RPM probes to monitor the BGP neighbors and determine if they are active.

Related Documentation

- [RPM Configuration Options on page 362](#)
- [RPM Support for VPN Routing and Forwarding on page 349](#)
- [Example: Configuring Basic RPM Probes on page 350](#)
- [Monitoring RPM Probes on page 366](#)
- [Determine What Causes Jitter and Latency on the Multilink Bundle on page 596](#)

IPv6 RPM Probes

Supported Platforms [vSRX](#)

Starting with Junos OS Release 15.1X49-D10, Route Engine-based RPM can send and receive IPv6 probe packets to monitor performance on IPv6 networks.

A probe request is a standard IPv6 packet with corresponding TCP, UDP, and ICMPv6 headers. A probe response is also a standard IPv6 packet with corresponding TCP, UDP, and ICMPv6 headers. No RPM header is appended to the standard packet for RE-based RPM. An IPv6-based RPM test occurs between an IPv6 RPM client and IPv6 RPM server.



NOTE: You can have both IPv4 tests and IPv6 tests in the same probe.

**Related
Documentation**

- [Guidelines for Configuring RPM Probes for IPv6 on page 348](#)
- [Configuring IPv6 RPM Probes on page 360](#)

Two-Way Active Measurement Protocol (TWAMP) Overview

Supported Platforms [SRX Series](#)

The Two-Way Active Measurement Protocol (TWAMP) is a standard protocol framework that defines the following separation of sessions based on the client/server architecture:

- **Control Connection**—The control connection negotiates the setting up of one or more test sessions between a TWAMP client and server. It is used to set up performance measurement sessions.


The remote operation process or daemon (rmopd) in the Junos OS Routing Engine (RE) takes care of all the control plane handling that includes both the client and server control message exchanges. The packet path processing for the control connection is the same as any other control process. Once the server-greeting, setup-response, and setup-start messages are successful, the TWAMP client starts the session creation message processing. One or more test sessions are created with the following tuples:

- Client/server Dynamic IP (DIP)
 - Client/server Source IP (SIP)
 - Client/server UDP port
- **Test Session**—The test session exchanges the probe packets to measure the performance metrics. Each test session between a client and server can use different QoS values to test QoS behavior on the network path for different packet priorities.

The timestamps on the test probes on the transmitter side should be just before the packet is sent out on the link or wire. Similarly, timestamps on the test probes on the receiver side should be done as soon as the probe packets are received from the link or wire. This will help in calculating a fairly accurate round trip time (RTT), delay, and jitter characteristics of the probes. The test probes can also be set to different QoS values to test the QoS behavior on the network path.

[Figure 8 on page 347](#) explains the TWAMP elements, which are implemented as client and server.

Figure 8: TWAMP Test




SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500 devices have the following limitations for TWAMP support:

- TWAMP for IPv6 is not supported.
- Authentication is not supported.

Four different TWAMP devices can perform the four logical roles of TWAMP control-client, server, session-sender, and session-reflector. A common implementation combines the roles of control-client and session-sender in one device (known as the TWAMP controller or TWAMP client) and the roles of server and session-reflector in the other device (known as the TWAMP responder or TWAMP server). In this case, each device runs both the TWAMP-Control (between control-client and server) and TWAMP-Test (between session-sender and session-reflector) protocols.

The four elements are shown in [Figure 9 on page 347](#).

Figure 9: Four Elements of TWAMP



Benefits of TWAMP

TWAMP configuration helps you to activate, test, monitor, and troubleshoot your network end-to-end without a dedicated test device. TWAMP defines an open protocol for measuring two-way or round-trip metrics with greater accuracy than other methods by using timestamps (processing delays can be factored as well). TWAMP is often used to check service-level agreement (SLA) compliance, and the TWAMP feature is often presented in that context.

Although TWAMP is a two-way measurement protocol, it is still possible to measure one-way packet loss; that is, separate loss values for the forward and backward directions. This is possible because the reflector places its own sequence number in the packet.

Related Documentation

- *Configuring TWAMP*

Guidelines for Configuring RPM Probes for IPv6

Supported Platforms [vSRX](#)

Starting with Junos OS Release 15.1X49-D10, you can configure RPM Probes for IPv6.

Keep the following guidelines in mind when you configure IPv6 addresses for RPM destinations or servers:

- IPv6 RPM uses ICMPv6 probe requests. You cannot configure ICMP or ICMP timestamp probe types.
- Only Routing Engine-based RPM is supported for IPv6 targets including VRF support, specification of the size of the data portion of ICMPv6 probes, data pattern, and traffic class.
- You can configure probes with a combination of IPv4 and IPv6 tests. However, an individual test must be either IPv4 or IPv6.
- Routing Engine-based RPM does not support hardware-based, or one-way hardware-based timestamping.
- We recommend that you include the **probe-limit** statement at the **[edit services rpm]** hierarchy level to set the limit on concurrent probes to 10. Higher concurrent probes can result in higher spikes.
- SNMP set operation is permitted only on ICMP probes and it is not supported for other probe types.
- The following table describes the IPv6 special address prefixes that you cannot configure in a probe.

IPv6 Address Type	IPv6 Address Prefix
Node-Scoped Unicast	::1/128 is the loopback address
	::/128 is the unspecified address
IPv4-Mapped Addresses	::FFFF:0:0/96
IPv4-Compatible Addresses	:<ipv4-address>/96
Link-Scoped Unicast	fe80::/10
Unique-Local	fc00::/7
Documentation Prefix	2001:db8::/32

IPv6 Address Type	IPv6 Address Prefix
6to4	2002::/16
6bone	5f00::/8
ORCHID	2001:10::/28
Teredo	2001::/32
Default Route	::/0
Multicast	ff00::/8

- In Routing Engine-based RPM, route-trip time (RTT) spikes might occur because of queuing delays, even with a single test.
- Since RPM might open TCP and UDP ports to communicate between the RPM server and RPM client, we recommend that you use firewalls and distributed denial-of-service (DDoS) attack filters to protect against security threats.

**Related
Documentation**

- [Configuring IPv6 RPM Probes on page 360](#)

RPM Support for VPN Routing and Forwarding

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

Real-time performance monitoring (RPM) is supported on all Juniper Network devices.

VRF in a Layer 3 VPN implementation allows multiple instances of a routing table to coexist within the same device at the same time. Because the routing instances are independent, the same or overlapping IPv4 or IPv6 addresses can be used without conflicting each other.

RPM ICMP and UDP probe with VPN routing and forwarding (VRF) has been improved. In previous releases, the RPM probes specified to a VRF table were not handled by the real-time forwarding process (FWDD-RT). In Junos OS Release 10.0, RPM probes specified to a VRF table are handled by the FWDD-RT, thereby providing more accurate results.

This feature supports RPM ICMP and UDP probes configured with routing instances of type VRF.

**Related
Documentation**

- [RPM Overview on page 341](#)
- [RPM Configuration Options on page 362](#)
- [Monitoring RPM Probes on page 366](#)

Example: Configuring Basic RPM Probes

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

This example shows how to configure basic RPM probes to measure performance between two network endpoints.

- [Requirements on page 350](#)
- [Overview on page 350](#)
- [Configuration on page 350](#)
- [Verification on page 353](#)

Requirements

Before you begin:

- Establish basic connectivity.
- Configure network interfaces. See *Interfaces Feature Guide for Security Devices*.

Overview

In this example, you configure basic probes for two RPM owners, customerA and customerB. You configure the RPM test as icmp-test for customerA with a test interval of 15 seconds and specify a probe type as icmp-ping-timestamp, a probe timestamp, and a target address as 192.178.16.5. You then configure the RPM thresholds and corresponding SNMP traps to catch ingress (inbound) times greater than 3000 microseconds.

Then you configure the RPM test as http-test for customerB with a test interval of 30 seconds and specify a probe type as http-get and a target URL as http://customerB.net. Finally, you configure RPM thresholds and corresponding SNMP traps as probe-failure and test-failure to catch three or more successive lost probes and total lost probes of 10.



NOTE: On SRX300, SRX320, SRX340, SRX1500 devices and vSRX instances, when you configure basic RPM probes, the following combination of the configuration parameters is not supported:

Source address and destination port and next-hop.

Configuring RPM probe with these parameters prevents sending out RPM probes to a specified probe target. We recommend you to configure either the source address or destination port and next-hop to configure RPM probe.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set services rpm probe customerA test icmp-test probe-interval 15
set services rpm probe customerA test icmp-test probe-type icmp-ping-timestamp
set services rpm probe customerA test icmp-test hardware-timestamp
set services rpm probe customerA test icmp-test target address 192.178.16.5
set services rpm probe customerA test icmp-test thresholds ingress-time 3000
set services rpm probe customerA test icmp-test traps ingress-time-exceeded
set services rpm probe customerB test http-test probe-interval 30
set services rpm probe customerB test http-test probe-type http-get
set services rpm probe customerB test http-test target url http://customerB.net
set services rpm probe customerB test http-test thresholds successive-loss 3
set services rpm probe customerB test http-test thresholds total-loss 10
set services rpm probe customerB test http-test traps probe-failure
set services rpm probe customerB test http-test traps test-failure
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure basic RPM probes:

1. Configure the RPM.

```
[edit]
user@host# edit services rpm
```

2. Configure the RPM owners.

```
[edit services rpm]
user@host# set probe customerA
user@host# set probe customerB
```

3. Configure the RPM test for customerA.

```
[edit services rpm]
user@host# edit probe customerA
user@host# set test icmp-test probe-interval 15
user@host# set test icmp-test probe-type icmp-ping-timestamp
```

4. Specify a probe timestamp and a target address.

```
[edit services rpm probe customerA]
user@host# set test icmp-test hardware-timestamp
user@host# set test icmp-test target address 192.178.16.5
```

5. Configure RPM thresholds and corresponding SNMP traps.

```
[edit services rpm probe customerA]
user@host# set test icmp-test thresholds ingress-time 3000
user@host# set test icmp-test traps ingress-time-exceeded
```

6. Configure the RPM test for customerB.

```
[edit]
user@host# edit services rpm probe customerB
user@host# set test http-test probe-interval 30
```

7. Specify a probe type and a target URL.

```
[edit services rpm probe customerB]
user@host# set test http-test probe-type http-get
user@host# set test http-test target url http://customerB.net
```

8. Configure RPM thresholds and corresponding SNMP traps.

```
[edit services rpm probe customerB]
user@host# set test http-test thresholds successive-loss 3
user@host# set test http-test thresholds total-loss 10
user@host# set test http-test traps probe-failure
user@host# set test http-test traps test-failure
```

Results From configuration mode, confirm your configuration by entering the **show services rpm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services rpm
probe customerA {
  test icmp-test {
    probe-type icmp-ping-timestamp;
    target address 192.0.2.2;
    probe-interval 15;
    thresholds {
      ingress-time 3000;
    }
    traps ingress-time-exceeded;
    hardware-timestamp;
  }
}
probe customerB {
  test http-test {
    probe-type http-get
    target url http://customerB.net;
    probe-interval 30;
    thresholds {
      successive-loss 3;
      total-loss 10;
    }
    traps [ probe-failure test-failure ];
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying RPM Services on page 353](#)
- [Verifying RPM Statistics on page 353](#)

Verifying RPM Services

Purpose Verify that the RPM configuration is within the expected values.

Action From configuration mode, enter the **show services rpm** command. The output shows the values that are configured for RPM on the device.

Verifying RPM Statistics

Purpose Verify that the RPM probes are functioning and that the RPM statistics are within expected values.

Action From configuration mode, enter the **show services rpm probe-results** command.

```
user@host> show services rpm probe-results
```

```

Owner: customerD, Test: icmp-test
Probe type: icmp-ping-timestamp
Minimum Rtt: 312 usec, Maximum Rtt: 385 usec, Average Rtt: 331 usec,
Jitter Rtt: 73 usec, Stddev Rtt: 27 usec
Minimum egress time: 0 usec, Maximum egress time: 0 usec,
Average egress time: 0 usec, Jitter egress time: 0 usec,
Stddev egress time: 0 usec
Minimum ingress time: 0 usec, Maximum ingress time: 0 usec,
Average ingress time: 0 usec, Jitter ingress time: 0 usec,
Stddev ingress time: 0 usec
Probes sent: 5, Probes received: 5, Loss percentage: 0

Owner: customerE, Test: http-test
Target address: 192.176.17.4, Target URL: http://customerB.net,
Probe type: http-get
Minimum Rtt: 1093 usec, Maximum Rtt: 1372 usec, Average Rtt: 1231 usec,
Jitter Rtt: 279 usec, Stddev Rtt: 114 usec
Probes sent: 3, Probes received: 3, Loss percentage: 0

Owner: Rpm-Bgp-Owner, Test: Rpm-Bgp-Test-1
Target address: 10.209.152.37, Probe type: icmp-ping, Test size: 5 probes
Routing Instance Name: LR1/RI1
Probe results:
  Response received, Fri Oct 28 05:20:23 2005
  Rtt: 662 usec
Results over current test:
  Probes sent: 5, Probes received: 5, Loss percentage: 0
  Measurement: Round trip time
    Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
    Jitter: 133 usec, Stddev: 53 usec
Results over all tests:
```

Probes sent: 5, Probes received: 5, Loss percentage: 0
Measurement: Round trip time
Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
Jitter: 133 usec, Stddev: 53 usec

- Related Documentation**
- [RPM Overview on page 341](#)
 - [RPM Configuration Options on page 362](#)
 - [Tuning RPM Probes on page 361](#)

Example: Configuring RPM Using TCP and UDP Probes

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

This example shows how to configure RPM using TCP and UDP probes.

- [Requirements on page 354](#)
- [Overview on page 354](#)
- [Configuration on page 354](#)
- [Verification on page 356](#)

Requirements

Before you begin:

- Establish basic connectivity.
- Configure network interfaces. See *Interfaces Feature Guide for Security Devices*.
- Configure the probe owner, the test, and the specific parameters of the RPM probe. See [“Example: Configuring Basic RPM Probes” on page 350](#).

Overview

In this example, you configure both the host (device A) and the remote device (device B) to act as TCP and UDP servers. You configure a probe for customerC, which uses TCP packets. Device B is configured as an RPM server for both TCP and UDP packets, using an lt services interface as the destination interface, and ports 50000 and 50037, respectively.



CAUTION: Use probe classification with caution, because improper configuration can cause packets to be dropped.

Configuration

- CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
{device A}
set services rpm probe customerC test tcp-test probe-interval 5
set services rpm probe customerC test tcp-test probe-type tcp-ping
set services rpm probe customerC test tcp-test target address 192.162.45.6
set services rpm probe customerC test tcp-test destination-interface lt-0/0/0
set services rpm probe customerC test tcp-test destination-port 50000

{device B}
set services rpm probe-server tcp port 50000
set services rpm probe-server udp port 50037
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure RPM using TCP and UDP probes:

1. Configure the RPM owner on device A.

```
{device A}
[edit]
user@host# edit services rpm
user@host# set probe customerC
```

2. Configure the RPM test.

```
{device A}
[edit services rpm]
user@host# edit services rpm probe customerC
user@host# set test tcp-test probe-interval 5
```

3. Set the probe type.

```
{device A}
[edit services rpm probe customerC]
user@host# set test tcp-test probe-type tcp-ping
```

4. Specify the target address.

```
{device A}
[edit services rpm probe customerC]
user@host# set test tcp-test target address 192.162.45.6
```

5. Configure the destination interface.

```
{device A}
[edit services rpm probe customerC]
user@host# set test tcp-test destination-interface lt-0/0/0
```

6. Configure port 50000 as the TCP port to which the RPM probes are sent.

```
{device A}
[edit services rpm probe customerC]
user@host# set test tcp-test destination-port 50000
```

7. Configure device B to act as a TCP server using port 50000.

```
{device B}
[edit]
user@host# edit services rpm
user@host# set probe-server tcp port 50000
```

8. Configure device B to act as a UDP server using port 50037.

```
{device B}
[edit services rpm]
user@host# set probe-server udp port 50037
```

Results From configuration mode, confirm your configuration by entering the **show services rpm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services rpm
  probe customerC {
    test tcp-test {
      probe-type tcp-ping;
      target address 192.162.45.6;
      probe-interval 5;
      destination-port 50000;
      destination-interface lt-0/0/0.0;
    }
  }
  probe-server {
    tcp {
      port 50000;
    }
    udp {
      port 50037;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying RPM Probe Servers

Purpose Confirm that the configuration is working properly.

Verify that the device is configured to receive and transmit TCP and UDP RPM probes on the correct ports.

Action From configuration mode, enter the **show services rpm active-servers** command. The output shows a list of the protocols and corresponding ports for which the device is configured as an RPM server.

```
user@host> show services rpm active-servers
```

```
Protocol: TCP, Port: 50000
```

```
Protocol: UDP, Port: 50037
```

- Related Documentation**
- [RPM Overview on page 341](#)
 - [RPM Configuration Options on page 362](#)
 - [Example: Configuring Basic RPM Probes on page 350](#)
 - [Example: Configuring RPM Probes for BGP Monitoring on page 357](#)
 - [Tuning RPM Probes on page 361](#)

Example: Configuring RPM Probes for BGP Monitoring

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

This example shows how to configure RPM probes to monitor BGP neighbors.

- [Requirements on page 357](#)
- [Overview on page 357](#)
- [Configuration on page 358](#)
- [Verification on page 359](#)

Requirements

Before you begin:

- Configure the BGP parameters under RPM configuration to send RPM probes to BGP neighbors. See [“Example: Configuring Basic RPM Probes” on page 350](#).
- Use TCP or UDP probes by configure both the probe server (Juniper Networks device) and the probe receiver (the remote device) to transmit and receive RPM probes on the same TCP or UDP port. See [“Example: Configuring RPM Using TCP and UDP Probes” on page 354](#).

Overview

In this example, you specify a hexadecimal value that you want to use for the data portion of the RPM probe as ABCD123. (It ranges from 1 through 2048 characters.) You specify the data size of the RPM probe as 1024 bytes. (The value ranges from 0 through 65,507.)

Then you configure destination port 50000 as the TCP port to which the RPM probes are sent. You specify the number of probe results to be saved in the probe history as 25.

(It ranges from 0 through 255, and the default is 50.) You set the probe count to 5 and probe interval as 1. (The probe count ranges from 1 through 15, and the default is 1; and the probe interval ranges from 1 through 255, and the default is 3.) You then specify tcp-ping as the type of probe to be sent as part of the test.

Finally, you set the test interval as 60. The value ranges from 0 through 86,400 seconds for the interval between tests.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set services rpm bgp data-fill ABCD123 data-size 1024
set services rpm bgp destination-port 50000 history-size 25
set services rpm bgp probe-count 5 probe-interval 1
set services rpm bgp probe-type tcp-ping test-interval 60
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure RPM probes to monitor BGP neighbors:

1. Configure the RPM and BGP.

```
[edit]
user@host# edit services rpm bgp
```
2. Specify a hexadecimal value.

```
[edit services rpm bgp]
user@host# set data-fill ABCD123
```
3. Specify the data size of the RPM probe.

```
[edit services rpm bgp]
user@host# set data-size 1024
```
4. Configure the destination port.

```
[edit services rpm bgp]
user@host# set destination-port 50000
```
5. Specify the number of probes.

```
[edit services rpm bgp]
user@host# set history-size 25
```
6. Set the probe count and probe interval.


```
[edit services rpm bgp]
user@host# set probe-count 5 probe-interval 1
```

7. Specify the type of probe.

```
[edit services rpm bgp]
user@host# set probe-type tcp-ping
```



NOTE: If you do not specify the probe type the default ICMP probes are sent.

8. Set the test interval.

```
[edit services rpm bgp]
user@host# set test-interval 60
```

Results From configuration mode, confirm your configuration by entering the **show services rpm** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show services rpm
bgp {
  probe-type tcp-ping;
  probe-count 5;
  probe-interval 1;
  test-interval 60;
  destination-port 50000;
  history-size 25;
  data-size 1024;
  data-fill ABCD123;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying RPM Probes for BGP Monitoring

Purpose Confirm that the configuration is working properly.

Verify that the RPM probes for BGP monitoring is configured.

Action From configuration mode, enter the **show services rpm** command.

Related Documentation

- [RPM Overview on page 341](#)
- [RPM Configuration Options on page 362](#)

- [Directing RPM Probes to Select BGP Devices on page 360](#)
- [Tuning RPM Probes on page 361](#)

Directing RPM Probes to Select BGP Devices

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

If a device has a large number of BGP neighbors configured, you can direct (filter) the RPM probes to a selected group of BGP neighbors rather than to all the neighbors. To identify the BGP devices to receive RPM probes, you can configure routing instances.

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To direct RPM probes to select BGP neighbors:

1. Configure routing instance **R11** to send RPM probes to BGP neighbors within the routing instance.

```
[edit services rpm bgp]  
user@host# set routing-instances R11
```

2. If you are done configuring the device, enter **commit** from configuration mode.

**Related
Documentation**

- [RPM Overview on page 341](#)
- [RPM Configuration Options on page 362](#)
- [Example: Configuring Basic RPM Probes on page 350](#)
- [Example: Configuring RPM Probes for BGP Monitoring on page 357](#)
- [Tuning RPM Probes on page 361](#)

Configuring IPv6 RPM Probes

Supported Platforms [vSRX](#)

Starting with Junos OS Release 15.1X49-D10, you can configure IPv6 destination addresses for an IPv6-based RPM probe test.

To configure an IPv6 RPM test:

1. Specify the RPM probe owner for the probe you want to configure as an IPv6 test.

```
[edit services rpm]  
user@host# edit probe customerA
```

2. Specify a name for the test.

```
[edit services rpm probe customerA]
```

```
user@host# edit test ipv6-test
```

3. Specify the probe type.

```
[edit services rpm probe customerA test ipv6-test]
user@host# set probe-type icmp6-ping
```

4. Specify the target address for the test.

```
[edit services rpm probe customerA test ipv6-test]
user@host# set target inet6-address 2001::2
```

5. Configure the remaining RPM test parameters.

Related Documentation

- [Guidelines for Configuring RPM Probes for IPv6 on page 348](#)

Tuning RPM Probes

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

After configuring an RPM probe, you can set parameters to control probe functions, such as the interval between probes, the total number of concurrent probes that a system can handle, and the source address used for each probe packet. See [“Example: Configuring Basic RPM Probes” on page 350](#).

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To tune RPM probes:

1. Set the maximum number of concurrent probes allowed on the system to 10.

```
[edit services rpm]
user@host# set probe-limit 10
```

2. Access the ICMP probe of customer A.

```
[edit]
user@host# edit services rpm probe customerA test icmp-test
```

3. Set the time between probe transmissions to 15 seconds.

```
[edit services rpm probe customerA test icmp-test]
user@host# set probe-interval 15
```

4. Set the number of probes within a test to 10.

```
[edit services rpm probe customerA test icmp-test]
user@host# set probe-count 10
```

5. Set the source address for each probe packet to **192.168.2.9**. If you do not explicitly configure a source address, the address on the outgoing interface through which the probe is sent is used as the source address.

```
[edit services rpm probe customerA test icmp-test]
user@host# set source-address 192.168.2.9
```

6. If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- [RPM Overview on page 341](#)
- [RPM Configuration Options on page 362](#)
- [Example: Configuring RPM Probes for BGP Monitoring on page 357](#)

RPM Configuration Options

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

You can configure real-time performance monitoring (RPM) parameters. See [Table 42 on page 362](#) for a summary of the configuration options.

Table 42: RPM Configuration Summary

Field	Function	Your Action
Performance Probe Owners		
Owner Name (required)	Identifies an RPM owner for which one or more RPM tests are configured. In most implementations, the owner name identifies a network on which a set of tests is being run (a particular customer, for example).	Type the name of the RPM owner.
Identification		
Test name (required)	Uniquely identifies the RPM test	Type the name of the RPM test.
Target (Address or URL) (required)	IPv4 or IPv6 address or URL of probe target	Type the IPv4 address, in dotted decimal notation, IPv6 address, or the URL of the probe target. If the target is a URL, type a fully formed URL that includes http:// .
Source Address	Explicitly configured IPv4 or IPv6 address to be used as the probe source address	Type the source address to be used for the probe. If the source address is not one of the device's assigned addresses, the packet uses the outgoing interface's address as its source.
Routing Instance	Particular routing instance over which the probe is sent	Type the routing instance name. The routing instance applies only to probes of type icmp , icmp6-ping , and icmp-timestamp . The default routing instance is inet.0 .
History Size	Number of probe results saved in the probe history	Type a number between 0 and 255. The default history size is 50 probes.

Table 42: RPM Configuration Summary (*continued*)

Field	Function	Your Action
Request Information		
Probe Type (required)	Specifies the type of probe to send as part of the test.	Select the desired probe type from the list: <ul style="list-style-type: none"> • <code>http-get</code> • <code>http-get-metadata</code> • <code>icmp6-ping</code> • <code>icmp-ping</code> • <code>icmp-ping-timestamp</code> • <code>tcp-ping</code> • <code>udp-ping</code>
Interval	Sets the wait time (in seconds) between each probe transmission	Type a number between 1 and 255 (seconds).
Test Interval (required)	Sets the wait time (in seconds) between tests.	Type a number between 0 and 86400 (seconds).
Probe Count	Sets the total number of probes to be sent for each test.	Type a number between 1 and 15.
Destination Port	Specifies the TCP or UDP port to which probes are sent. To use TCP or UDP probes, you must configure the remote server as a probe receiver. Both the probe server and the remote server must be Juniper Networks devices configured to receive and transmit RPM probes on the same TCP or UDP port.	Type the number 7—a standard TCP or UDP port number—or a port number from 49152 through 65535.
DSCP Bits	Specifies the Differentiated Services code point (DSCP) bits. This value must be a valid 6-bit pattern. The default is 000000 .	Type a valid 6-bit pattern.
Data Size	Specifies the size of the data portion of the ICMP probes.	Type a size (in bytes) between 0 and 65507.
Data Fill	Specifies the contents of the data portion of the ICMP probes.	Type a hexadecimal value between 1 and 800h to use as the contents of the ICMP probe data.
Hardware Timestamp	Enables timestamping of RPM probe messages. You can timestamp the following RPM probes to improve the measurement of latency or jitter: <ul style="list-style-type: none"> • ICMP ping • ICMP ping timestamp • UDP ping—destination port UDP-ECHO (port 7) only • UDP ping timestamp—destination port UDP-ECHO (port 7) only 	To enable timestamping, select the check box.

Table 42: RPM Configuration Summary (*continued*)

Field	Function	Your Action
Maximum Probe Thresholds		
Successive Lost Probes	Sets the total number of probes that must be lost successively to trigger a probe failure and generate a system log message.	Type a number between 0 and 15.
Lost Probes	Sets the total number of probes that must be lost to trigger a probe failure and generate a system log message.	Type a number between 0 and 15.
Round Trip Time	Sets the total round-trip time (in microseconds), from the device to the remote server, that triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Jitter	Sets the total jitter (in microseconds), for a test, that triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Standard Deviation	Sets the maximum allowable standard deviation (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Egress Time	Sets the total one-way time (in microseconds), from the device to the remote server, that triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Ingress Time	Sets the total one-way time (in microseconds), from the remote server to the device, that triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds)
Jitter Egress Time	Sets the total outbound-time jitter (in microseconds), for a test, that triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds)
Jitter Ingress Time	Sets the total inbound-time jitter (in microseconds), for a test, that triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Egress Standard Deviation	Sets the maximum allowable standard deviation of outbound times (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Ingress Standard Deviation	Sets the maximum allowable standard deviation of inbound times (in microseconds) for a test, which, if exceeded, triggers a probe failure and generates a system log message.	Type a number between 0 and 60,000,000 (microseconds).
Traps		

Table 42: RPM Configuration Summary (*continued*)

Field	Function	Your Action
Egress Jitter Exceeded	Generates SNMP traps when the threshold for jitter in outbound time is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Egress Standard Deviation Exceeded	Generates SNMP traps when the threshold for standard deviation in outbound times is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Egress Time Exceeded	Generates SNMP traps when the threshold for maximum outbound time is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Ingress Jitter Exceeded	Generates SNMP traps when the threshold for jitter in inbound time is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Ingress Standard Deviation Exceeded	Generates SNMP traps when the threshold for standard deviation in inbound times is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Ingress Time Exceeded	Generates traps when the threshold for maximum inbound time is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Jitter Exceeded	Generates traps when the threshold for jitter in round-trip time is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Probe Failure	Generates traps when the threshold for the number of successive lost probes is reached.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
RTT Exceeded	Generates traps when the threshold for maximum round-trip time is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Standard Deviation Exceeded	Generates traps when the threshold for standard deviation in round-trip times is exceeded.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Test Completion	Generates traps when a test is completed.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.
Test Failure	Generates traps when the threshold for the total number of lost probes is reached.	<ul style="list-style-type: none"> To enable SNMP traps for this condition, select the check box. To disable SNMP traps, clear the check box.

Performance Probe Server

Table 42: RPM Configuration Summary (*continued*)

Field	Function	Your Action
TCP Probe Server	Specifies the port on which the device is to receive and transmit TCP probes.	Type the number 7—a standard TCP or UDP port number—or a port number from 49160 through 65535.
UDP Probe Server	Specifies the port on which the device is to receive and transmit UDP probes.	Type the number 7—a standard TCP or UDP port number—or a port number from 49160 through 65535.



NOTE: On SRX300, SRX320, SRX340, SRX1500 devices and vSRX instances, when you configure basic RPM probes, the following combination of the configuration parameters is not supported:

Source address and destination port and next-hop.

Configuring RPM probe with these parameters prevents sending out RPM probes to a specified probe target. We recommend you to configure either the source address or destination port and next-hop to configure RPM probe.

Related Documentation

- [RPM Overview on page 341](#)
- [Example: Configuring Basic RPM Probes on page 350](#)
- [Example: Configuring RPM Using TCP and UDP Probes on page 354](#)
- [Example: Configuring RPM Probes for BGP Monitoring on page 357](#)

Monitoring RPM Probes


Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

The RPM information includes the round-trip time, jitter, and standard deviation values for each configured RPM test on the device. To view these RPM properties, select **Troubleshoot > RPM > View RPM** in the J-Web user interface, or in configuration mode enter the **show** command:

```
[edit]
user@host# run show services rpm probe-results
```

In addition to the RPM statistics for each RPM test, the J-Web user interface displays the round-trip times and cumulative jitter graphically. [Figure 10 on page 367](#) shows sample graphs for an RPM test.

Figure 10: Sample RPM Graphs



In Figure 10 on page 367, the round-trip time and jitter values are plotted as a function of the system time. Large spikes in round-trip time or jitter indicate a slower outbound (egress) or inbound (ingress) time for the probe sent at that particular time.

Table 43 on page 367 summarizes key output fields in RPM displays.

Table 43: Summary of Key RPM Output Fields

Field	Values	Additional Information
Currently Running Tests		
Graph		Click the Graph link to display the graph (if it is not already displayed) or to update the graph for a particular test.
Owner	Configured owner name of the RPM test.	—
Test Name	Configured name of the RPM test.	—

Table 43: Summary of Key RPM Output Fields (*continued*)

Field	Values	Additional Information
Probe Type	Type of RPM probe configured for the specified test: <ul style="list-style-type: none"> • http-get • http-get-metadata • icmp-ping • icmp6-ping • icmp-ping-timestamp • tcp-ping • udp-ping 	—
Target Address	IPv4 address, IPv6 address, or URL of the remote server that is being probed by the RPM test.	—
Source Address	Explicitly configured IPv4 or IPv6 source address that is included in the probe packet headers.	If no source address is configured, the RPM probe packets use the outgoing interface as the source address, and the Source Address field is empty.
Minimum RTT	Shortest round-trip time from the Juniper Networks device to the remote server, as measured over the course of the test.	—
Maximum RTT	Longest round-trip time from the Juniper Networks device to the remote server, as measured over the course of the test.	—
Average RTT	Average round-trip time from the Juniper Networks device to the remote server, as measured over the course of the test.	—
Standard Deviation RTT	Standard deviation of round-trip times from the Juniper Networks device to the remote server, as measured over the course of the test.	—
Probes Sent	Total number of probes sent over the course of the test.	—
Loss Percentage	Percentage of probes sent for which a response was not received.	—
Round-Trip Time for a Probe		
Samples	Total number of probes used for the data set.	The Juniper Networks device maintains records of the most recent 50 probes for each configured test. These 50 probes are used to generate RPM statistics for a particular test.
Earliest Sample	System time when the first probe in the sample was received.	—
Latest Sample	System time when the last probe in the sample was received.	—

Table 43: Summary of Key RPM Output Fields (*continued*)

Field	Values	Additional Information
Mean Value	Average round-trip time for the 50-probe sample.	–
Standard Deviation	Standard deviation of the round-trip times for the 50-probe sample.	–
Lowest Value	Shortest round-trip time from the device to the remote server, as measured over the 50-probe sample.	–
Time of Lowest Sample	System time when the lowest value in the 50-probe sample was received.	–
Highest Value	Longest round-trip time from the Juniper Networks device to the remote server, as measured over the 50-probe sample.	–
Time of Highest Sample	System time when the highest value in the 50-probe sample was received.	–
Cumulative Jitter for a Probe		
Samples	Total number of probes used for the data set.	The Juniper Networks device maintains records of the most recent 50 probes for each configured test. These 50 probes are used to generate RPM statistics for a particular test.
Earliest Sample	System time when the first probe in the sample was received.	–
Latest Sample	System time when the last probe in the sample was received.	–
Mean Value	Average jitter for the 50-probe sample.	–
Standard Deviation	Standard deviation of the jitter values for the 50-probe sample.	–
Lowest Value	Smallest jitter value, as measured over the 50-probe sample.	–
Time of Lowest Sample	System time when the lowest value in the 50-probe sample was received.	–
Highest Value	Highest jitter value, as measured over the 50-probe sample.	–
Time of Highest Sample	System time when the highest jitter value in the 50-probe sample was received.	–

- Related Documentation**
- [RPM Overview on page 341](#)
 - [RPM Support for VPN Routing and Forwarding on page 349](#)
 - [RPM Configuration Options on page 362](#)

CHAPTER 21

Configuring IP Monitoring

- [IP Monitoring Overview on page 371](#)
- [Understanding IP Monitoring Test Parameters on page 372](#)
- [Example: Configuring IP Monitoring on SRX Series Devices on page 373](#)
- [Understanding IP Monitoring Through Redundant Ethernet Interface Link Aggregation Groups on page 376](#)
- [Example: Configuring IP Monitoring on SRX Series Devices on page 376](#)
- [Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring on page 382](#)

IP Monitoring Overview

Supported Platforms [SRX Series, vSRX](#)

This feature monitors IP on standalone SRX Series devices or a chassis cluster redundant Ethernet (reth) interface. Existing RPM probes are sent to an IP address to check for reachability. The user takes action based on the reachability result. Supported action currently is preferred static route injection to system route table.

The actions supported are:

- Adding or deleting a new static route that has a higher priority (lower preference) value than a route configured through the CLI command **set routing-options static route**
- Defining multiple probe names under the same IP monitoring policy. If any probe fails, the action is taken. If all probes are reachable, the action is reverted
- Configuring multiple tests in one RPM probe. All tests must fail for the RPM probe to be considered unreachable. If at least one test reaches its target, the RPM probe is considered reachable
- Configuring multiple failure thresholds in one RPM test. If one threshold is reached, the test fails. If no thresholds are reached, the test succeeds.
- Specifying the no-preempt option. If the no-preempt option is specified, the policy does not perform preemptive failback when it is in a failover state or when the RPM probe test recovers from a failure.

- Setting preferred metric values. If the preferred metric value is set, during failover, the route is injected with the set preferred metric value.
- Enabling and disabling interfaces.
 - **Interface-Enable**—On a physical or logical interface, when the interface-enable action is configured, the initial state of the interface is disable after startup, and it continues to remain in the disable state as long as the associated RPM probe is in the pass state. When the associated RPM probe fails, the configured physical and logical interfaces are enabled.
 - **Interface-Disable**—On a physical or logical interface, when the interface-disable action is configured, the interface state remains unchanged. When the associated RPM probe fails, the physical and logical interfaces are disabled.



NOTE: Multiple probe names and actions can be defined for the same IP monitoring policy.

Related Documentation

- [Understanding IP Monitoring Test Parameters on page 372](#)

Understanding IP Monitoring Test Parameters

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

Each probed target is monitored over the course of a test, which represents a collection of probes during which statistics such as standard deviation and jitter are collected are calculated. During a test, probes are generated and responses collected at a rate defined by the probe interval, the number of seconds between probes.



NOTE: To avoid flap, an action is reverted only at the end of a test cycle. During the test cycle, if no threshold is reached, the action is reverted. Although action-failover takes place based on a predefined condition of a monitored IP, when the condition is reversed, the IP becomes reachable on the original route, and the newly added route is deleted. Recovery is performed only when all RPM probes report the IP as reachable.

[Table 44 on page 372](#) lists the test parameters and its default values:

Table 44: Test Parameters and Default Values

Parameter	Default Value
probe-count	1
probe-interval	3 seconds
test-interval	1 second

Table 45 on page 373 lists the supported threshold and its description:

Table 45: Threshold Supported and Description

Threshold	Description
Successive-Loss	Successive loss count of probes
Total-Loss	Total probe lost count

Related Documentation

- [IP Monitoring Overview on page 371](#)

Example: Configuring IP Monitoring on SRX Series Devices

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX

This example shows how to monitor IP on SRX300, SRX320, SRX340, SRX550M, and SRX1500 devices.

- [Requirements on page 373](#)
- [Overview on page 373](#)
- [Configuration on page 373](#)
- [Verification on page 375](#)

Requirements

Before you begin:

Configure the following RPM options for RPM test:

- target-address
- probe-count
- probe-interval
- test-interval
- thresholds
- next-hop

Overview

This example shows how to set up IP monitoring on an SRX Series device.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set services rpm probe Probe-Payment-Server test paysvr target address 1.1.1.10
set services rpm probe Probe-Payment-Server test paysvr probe-count 10
set services rpm probe Probe-Payment-Server test paysvr probe-interval 5
set services rpm probe Probe-Payment-Server test paysvr test-interval 5
set services rpm probe Probe-Payment-Server test paysvr thresholds successive-loss 10
set services rpm probe Probe-Payment-Server test paysvr next-hop 2.2.2.1
set services ip-monitoring policy Payment-Server-Tracking match rpm-probe
Probe-Payment-Server
set services ip-monitoring policy Payment-Server-Tracking then preferred-route route
1.1.1.0/24 next-hop 1.1.1.99
```

**Step-by-Step
Procedure**

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IP monitoring on an SRX Series Services Gateway:

1. Configure the target address under the RPM probe.

```
[edit ]
user@host# set services rpm probe Probe-Payment-Server test paysvr target address
1.1.1.10
```

2. Configure the probe count under the RPM probe.

```
[edit ]
user@host# set services rpm probe Probe-Payment-Server test paysvr probe-count
10
```

3. Configure the probe interval (in seconds) under the RPM probe.

```
[edit ]
user@host# set services rpm probe Probe-Payment-Server test paysvr probe-interval
5
```

4. Configure the test interval (in seconds) under the RPM probe.

```
[edit ]
user@host# set services rpm probe Probe-Payment-Server test paysvr test-interval
5
```

5. Configure the threshold successive loss count under the RPM

```
[edit ]
user@host# set services rpm probe Probe-Payment-Server test paysvr thresholds
successive-loss 10
```

6. Configure the next-hop IP address under the RPM probe.

```
[edit ]
user@host# set services rpm probe Probe-Payment-Server test paysvr next-hop
2.2.2.1
```


7. Configure the IP monitoring policy under services.

```
[edit ]
user@host# set services ip-monitoring policy Payment-Server-Tracking match
rpm-probe Probe-Payment-Server
```



NOTE: The following steps are not mandatory. You can configure interface actions and route actions independently, or you can configure both the interface action and the route action together in one IP monitoring policy.

8. Configure the IP monitoring preferred route under services.

```
[edit ]
user@host# set services ip-monitoring policy Payment-Server-Tracking then
preferred-route route 1.1.1.0/24 preferred-metric 4
```

9. Configure the IP monitoring interface actions.

- Enable

```
[edit ]
user@host# set services ip-monitoring policy Payment-Server-Tracking then
interface ge-0/0/1 enable
```

- Disable

```
[edit ]
user@host# set services ip-monitoring policy Payment-Server-Tracking then
interface fe-0/0/[4-6] disable
```

10. Configure the no-preempt option.

```
[edit ]
user@host# set services ip-monitoring policy Payment-Server-Tracking no-preempt
```

Verification

Verifying IP Monitoring

Purpose Verify the IP monitoring status of a policy.

Action To verify the configuration is working properly, enter the following command:

```
show services ip-monitoring status <policy-name>
```

Related Documentation

- [IP Monitoring Overview on page 371](#)
- [Understanding IP Monitoring Test Parameters on page 372](#)

Understanding IP Monitoring Through Redundant Ethernet Interface Link Aggregation Groups

Supported Platforms [SRX1500, SRX5400, SRX5600, SRX5800, vSRX](#)

IP monitoring checks the reachability of an upstream device. It is designed to check the end-to-end connectivity of configured IP addresses and allows a redundancy group (RG) to automatically failover when the monitored IP address is not reachable through the redundant Ethernet. Both the primary and secondary devices in the chassis cluster monitor specific IP addresses to determine whether an upstream device in the network is reachable.

A redundant Ethernet interface contains physical interfaces from both the primary and secondary nodes in the SRX Series chassis cluster. In a redundant Ethernet interface, two physical interfaces are configured with each node contributing one physical interface. In a redundant Ethernet interface LAG, more than two physical interfaces are configured in the redundant Ethernet interface.

Related Documentation

- [IP Monitoring Overview on page 371](#)

Example: Configuring IP Monitoring on SRX Series Devices

Supported Platforms [SRX1500, SRX5600, SRX5800, vSRX](#)

This example shows how to monitor IP on SRX1500, SRX5600, and SRX5800 devices with chassis cluster enabled.

- [Requirements on page 376](#)
- [Overview on page 376](#)
- [Configuration on page 377](#)
- [Verification on page 379](#)

Requirements

- You need two SRX5800 Services Gateways with identical hardware configurations, one SRX Series device and one EX8208 Ethernet Switch.
- Physically connect the two SRX5800 devices (back-to-back for the fabric and control ports) and ensure that they are the same models. Configure/add these two devices in a cluster.

Overview

IP address monitoring checks end-to-end reachability of configured IP address and allows a redundancy group to automatically fail over when not reachable through the child link of redundant Ethernet interface (known as a reth) interface. Redundancy groups on both devices in a cluster can be configured to monitor specific IP addresses to determine whether an upstream device in the network is reachable.

This example shows how to set up IP monitoring on an SRX Series device.



NOTE: IP monitoring is not supported on an NP-IOC card.




NOTE: IP monitoring does not support MIC online/offline status on SRX devices.

Topology

Figure 11 on page 377 shows the topology used in this example.

Figure 11: IP Monitoring on an SRX Series Device Topology Example



In this example, two SRX5800 devices in a chassis cluster are connected to an SRX1500 device through an EX8208 Ethernet Switch. The example shows how the redundancy groups can be configured to monitor key upstream resources reachable through redundant Ethernet interfaces on either node in a cluster.

Configuration

- [Configuring IP Monitoring on SRX Series Device on page 378](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set chassis cluster reth-count 1
set chassis cluster redundancy-group 0 node 0 priority 254
set chassis cluster redundancy-group 0 node 1 priority 1
```

```
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 199
set chassis cluster redundancy-group 1 ip-monitoring global-weight 255
set chassis cluster redundancy-group 1 ip-monitoring global-threshold 80
set chassis cluster redundancy-group 1 ip-monitoring retry-interval 3
set chassis cluster redundancy-group 1 ip-monitoring retry-count 10
set chassis cluster redundancy-group 1 ip-monitoring family inet 192.0.0.1 weight
80
set chassis cluster redundancy-group 1 ip-monitoring family inet 192.0.0.1
interface reth0.0 secondary-ip-address 192.0.2.2
set interfaces ge-0/0/1 gigether-options redundant-parent reth0
set interfaces ge-4/0/1 gigether-options redundant-parent reth0
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 192.0.2.1/24
set routing-options static route 192.0.0.1/32 next-hop 192.0.2.3
```

Configuring IP Monitoring on SRX Series Device

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure IP monitoring on an SRX Series device:

1. Specify the number of redundant Ethernet interfaces.

```
{primary:node0}[edit]
user@host# set chassis cluster reth-count 1
```

2. Specify a redundancy group's priority for primacy on each node of the cluster. The higher number takes precedence.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 0 node 0 priority 254
user@host# set chassis cluster redundancy-group 0 node 1 priority 1
user@host# set chassis cluster redundancy-group 1 node 0 priority 200
user@host# set chassis cluster redundancy-group 1 node 1 priority 199
```

3. Configure the redundant Ethernet interfaces to redundancy-group 1.

```
{primary:node0}[edit]
user@host# set interfaces reth0 redundant-ether-options redundancy-group 1
user@host# set interfaces reth0 unit 0 family inet address 192.0.2.1/24
```

4. Assign child interfaces for the redundant Ethernet interfaces from node 0 and node 1.

```
{primary:node0}[edit]
user@host# set interfaces ge-0/0/1 gigether-options redundant-parent reth0
user@host# set interfaces ge-4/0/1 gigether-options redundant-parent reth0
```

5. Configure the static route to the IP address that is to be monitored.

```
{primary:node0}[edit]
user@host# set routing-options static route 192.0.0.1/32 next-hop 192.0.2.3
```

6. Configure IP monitoring under redundancy-group 1 with global weight and global threshold.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring global-weight
255
user@host# set chassis cluster redundancy-group 1 ip-monitoring global-threshold
80
```

7. Specify the retry interval.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring retry-interval 3
```

8. Specify the retry count.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring retry-count 10
```

9. Assign a weight to the IP address to be monitored, and configure a secondary IP address that will be used to send ICMP packets from the secondary node to track the IP being monitored.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring family inet 192.0.0.1
weight 80
user@host# set chassis cluster redundancy-group 1 ip-monitoring family inet 192.0.0.1
interface reth0.0 secondary-ip-address 192.0.2.2
```



NOTE:

- The redundant Ethernet (reth0) IP address, 192.0.2.1/24, is used to send ICMP packets from node 0 to check the reachability of the monitored IP.
- The secondary IP address, 192.0.2.2, should belong to the same network as the reth0 IP address.
- The secondary IP address is used to send ICMP packets from node 1 to check the reachability of the monitored IP.

Verification

Confirm the configuration is working properly.

- [Verifying Chassis Cluster Status—Before Failover on page 380](#)
- [Verifying Chassis Cluster IP Monitoring Status—Before Failover on page 380](#)
- [Verifying Chassis Cluster Status—After Failover on page 380](#)
- [Verifying Chassis Cluster IP Monitoring Status—After Failover on page 381](#)

Verifying Chassis Cluster Status—Before Failover

Purpose Verify the chassis cluster status, failover status, and redundancy group information before failover.

Action From operational mode, enter the **show chassis cluster status** command.

```
show chassis cluster status
```

```
Cluster ID: 11
```

Node	Priority	Status	Preempt	Manual failover
------	----------	--------	---------	-----------------

```
Redundancy group: 0 , Failover count: 0
```

node0	254	primary	no	no
node1	1	secondary	no	no

```
Redundancy group: 1 , Failover count: 0
```

node0	200	primary	no	no
node1	199	secondary	no	no

Verifying Chassis Cluster IP Monitoring Status—Before Failover

Purpose Verify the IP status being monitored from both nodes and the failover count for both nodes before failover.

Action From operational mode, enter the **show chassis cluster ip-monitoring status redundancy-group 1** command.

```
show chassis cluster ip-monitoring status redundancy-group 1
```

```
node0:
```

```
-----
```

```
Redundancy group: 1
```

IP address	Status	Failure count	Reason
192.0.0.1	reachable	0	n/a

```
node1:
```

```
-----
```

```
Redundancy group: 1
```

IP address	Status	Failure count	Reason
192.0.0.1	reachable	0	n/a

Verifying Chassis Cluster Status—After Failover

Purpose Verify the chassis cluster status, failover status, and redundancy group information after failover.



NOTE: If the IP address is not reachable, the following output will be displayed.

Action From operational mode, enter the **show chassis cluster status** command.

```
show chassis cluster status
```

Cluster ID: 11

Node	Priority	Status	Preempt	Manual failover
------	----------	--------	---------	-----------------

Redundancy group: 0 , Failover count: 0

node0	254	primary	no	no
node1	1	secondary	no	no

Redundancy group: 1 , Failover count: 1

node0	0	secondary	no	no
node1	199	primary	no	no

Verifying Chassis Cluster IP Monitoring Status—After Failover

Purpose Verify the IP status being monitored from both nodes and the failover count for both nodes after failover.

Action From operational mode, enter the **show chassis cluster ip-monitoring status redundancy-group 1** command.

```
show chassis cluster ip-monitoring status redundancy-group 1
```

node0:

Redundancy group: 1

IP address	Status	Failure count	Reason
192.0.0.1	unreachable	1	unknown

node1:

Redundancy group: 1

IP address	Status	Failure count	Reason
192.0.0.1	reachable	0	n/a

Related Documentation • *Example: Configuring an Active/Passive Chassis Cluster on SRX Series Services Gateways*

Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring

Supported Platforms SRX Series, vSRX

This example shows how to configure redundancy group IP address monitoring for an SRX Series device in a chassis cluster.

- [Requirements on page 382](#)
- [Overview on page 382](#)
- [Configuration on page 383](#)
- [Verification on page 384](#)

Requirements

Before you begin:

- Set the chassis cluster node ID and cluster ID. See *Example: Setting the Chassis Cluster Node ID and Cluster ID for SRX Series Devices*.
- Configure the chassis cluster management interface. See *Example: Configuring the Chassis Cluster Management Interface*.
- Configure the chassis cluster fabric. See *Example: Configuring the Chassis Cluster Fabric Interfaces*.

Overview

You can configure redundancy groups to monitor upstream resources by pinging specific IP addresses that are reachable through redundant Ethernet interfaces on either node in a cluster. You can also configure global threshold, weight, retry interval, and retry count parameters for a redundancy group. When a monitored IP address becomes unreachable, the weight of that monitored IP address is deducted from the redundancy group IP address monitoring global threshold. When the global threshold reaches 0, the global weight is deducted from the redundancy group threshold. The retry interval determines the ping interval for each IP address monitored by the redundancy group. The pings are sent as soon as the configuration is committed. The retry count sets the number of allowed consecutive ping failures for each IP address monitored by the redundancy group.

In this example, you configure the following settings for redundancy group 1:

- IP address to monitor—10.1.1.10
- IP address monitoring global-weight—100
- IP address monitoring global-threshold—200



NOTE: The threshold applies cumulatively to all IP addresses monitored by the redundancy group.

- IP address retry-interval—3 seconds

- IP address retry-count—10
- Weight—150
- Redundant Ethernet interface—reth1.0
- Secondary IP address—10.1.1.101

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
{primary:node0}[edit]
user@host#
set chassis cluster redundancy-group 1 ip-monitoring global-weight 100
set chassis cluster redundancy-group 1 ip-monitoring global-threshold 200
set chassis cluster redundancy-group 1 ip-monitoring retry-interval 3
set chassis cluster redundancy-group 1 ip-monitoring retry-count 10
set chassis cluster redundancy-group 1 ip-monitoring family inet 10.1.1.10 weight 150
interface reth1.0 secondary-ip-address 10.1.1.101
```

Step-by-Step Procedure To configure redundancy group IP address monitoring:

1. Specify a global monitoring weight.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring global-weight
100
```

2. Specify the global monitoring threshold.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring global-threshold
200
```

3. Specify the retry interval.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring retry-interval 3
```

4. Specify the retry count.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring retry-count 10
```

5. Specify the IP address to be monitored, weight, redundant Ethernet interface, and secondary IP address.

```
{primary:node0}[edit]
user@host# set chassis cluster redundancy-group 1 ip-monitoring family inet 10.1.1.10
weight 100 interface reth1.0 secondary-ip-address 10.1.1.101
```

Results From configuration mode, confirm your configuration by entering the **show chassis cluster redundancy-group 1** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
{primary:node0}[edit]
user@host# show chassis cluster redundancy-group 1
ip-monitoring {
  global-weight 100;
  global-threshold 200;
  family {
    inet {
      10.1.1.10 {
        weight 100;
        interface reth1.0 secondary-ip-address 10.1.1.101;
      }
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Status of Monitored IP Addresses for a Redundancy Group

Purpose Verify the status of monitored IP addresses for a redundancy group.

Action From operational mode, enter the **show chassis cluster ip-monitoring status** command. For information about a specific group, enter the **show chassis cluster ip-monitoring status redundancy-group** command.

```
{primary:node0}
user@host> show chassis cluster ip-monitoring status
node0:
```

```
-----
Redundancy group: 1
Global threshold: 200
Current threshold: -120
```

IP address	Status	Failure count	Reason	Weight
10.1.1.10	reachable	0	n/a	220
10.1.1.101	reachable	0	n/a	100

```
node1:
```

```
-----
Redundancy group: 1
Global threshold: 200
Current threshold: -120
```

IP address	Status	Failure count	Reason	Weight
10.1.1.10	reachable	0	n/a	220
10.1.1.101	reachable	0	n/a	100

**Related
Documentation**

- *Understanding Chassis Cluster Redundancy Group Interface Monitoring*
- *Understanding Chassis Cluster Redundancy Group IP Address Monitoring*
- *Understanding Chassis Cluster Redundancy Group Failover*

PART 7

Monitoring Common Security Features

- [Displaying Real-Time Information from Device to Host on page 389](#)
- [Monitoring Application Layer Gateways Features on page 395](#)
- [Monitoring Interfaces and Switching Functions on page 421](#)
- [Monitoring NAT on page 441](#)
- [Monitoring Security Policies on page 453](#)
- [Monitoring Events, Services and System on page 479](#)
- [Monitoring Unified Threat Management Features on page 489](#)
- [Monitoring VPNs on page 503](#)

Displaying Real-Time Information from Device to Host

- [Displaying Real-Time Monitoring Information on page 389](#)
- [Displaying Multicast Path Information on page 391](#)

Displaying Real-Time Monitoring Information

Supported Platforms [SRX Series, vSRX](#)

To display real-time monitoring information about each device between the device and a specified destination host, enter the **traceroute monitor** command with the following syntax:

```
user@host> traceroute monitor host <count number> <inet | inet6> <interval seconds>
<no-resolve> <size bytes><source source-address> <summary>
```

[Table 46 on page 389](#) describes the **traceroute monitor** command options.

Table 46: CLI traceroute monitor Command Options

Option	Description
<i>host</i>	Sends traceroute packets to the hostname or IP address you specify.
<i>count number</i>	(Optional) Limits the number of ping requests, in packets, to send in summary mode. If you do not specify a count, ping requests are continuously sent until you press Q.
<i>inet</i>	(Optional) Forces the traceroute packets to an IPv4 destination.
<i>inet6</i>	(Optional) Forces the traceroute packets to an IPv6 destination.
<i>interval seconds</i>	(Optional) Sets the interval between ping requests, in seconds. The default value is 1 second.
<i>no-resolve</i>	(Optional) Suppresses the display of the hostnames of the hops along the path.
<i>size bytes</i>	(Optional) Sets the size of the ping request packet. The size can be from 0 through 65,468 bytes. The default packet size is 64 bytes.
<i>source address</i>	(Optional) Uses the source address that you specify, in the traceroute packet.

Table 46: CLI traceroute monitor Command Options (*continued*)

Option	Description
summary	(Optional) Displays the summary traceroute information.

To quit the **traceroute monitor** command, press Q.

The following is sample output from a **traceroute monitor** command:

```

user@host> traceroute monitor host2

My traceroute  [v0.69]
host (0.0.0.0)(tos=0x0 psize=64 bitpattern=0x00)
Wed Mar 14 23:14:11 2007
Keys:  Help   Display mode   Restart statistics   Order of fields   quit

Pings
Host
Last  Avg  Best  Wrst  StDev
1. 173.24.232.66      0.0%    5
9.4  8.6  4.8   9.9   2.1
2. 173.24.232.66      0.0%    5
7.9 17.2  7.9  29.4  11.0
3. 173.24.232.66      0.0%    5
9.9  9.3  8.7   9.9   0.5
4. 173.24.232.66      0.0%    5
9.9  9.8  9.5  10.0   0.2

```

Table 47 on page 390 summarizes the output fields of the display.

Table 47: CLI traceroute monitor Command Output Summary

Field	Description
host	Hostname or IP address of the device issuing the traceroute monitor command.
psizesize	Size of ping request packet, in bytes.
Keys	
Help	Displays the Help for the CLI commands. Press H to display the Help.
Display mode	Toggles the display mode. Press D to toggle the display mode
Restart statistics	Restarts the traceroute monitor command. Press R to restart the traceroute monitor command.
Order of fields	Sets the order of the displayed fields. Press O to set the order of the displayed fields.

Table 47: CLI traceroute monitor Command Output Summary (*continued*)

Field	Description
quit	Quits the traceroute monitor command. Press Q to quit the traceroute monitor command.
Packets	
<i>number</i>	Number of the hop (device) along the route to the final destination host.
Host	Hostname or IP address of the device at each hop.
Loss%	Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage.
Pings	
Snt	Number of ping requests sent to the device at this hop.
Last	Most recent round-trip time, in milliseconds, to the device at this hop.
Avg	Average round-trip time, in milliseconds, to the device at this hop.
Best	Shortest round-trip time, in milliseconds, to the device at this hop.
Wrst	Longest round-trip time, in milliseconds, to the device at this hop.
StDev	Standard deviation of round-trip times, in milliseconds, to the device at this hop.

Related Documentation

- [Displaying Log and Trace Files on page 538](#)

Displaying Multicast Path Information

Supported Platforms [SRX Series](#)

To display information about a multicast path from a source to the device, enter the **mtrace from-source** command with the following syntax:

```
user@host> mtrace from-source source host <extra-hops number> <group address>
<interval seconds> <max-hops number> <max-queries number> <response host>
<routing-instance routing-instance-name> <ttr number> <wait-time seconds> <loop>
<multicast-response | unicast-response> <no-resolve> <no-router-alert> <brief |
detail>
```

[Table 48 on page 391](#) describes the **mtrace from-source** command options.

Table 48: CLI mtrace from-source Command Options

Option	Description
source host	Traces the path to the specified hostname or IP address.

Table 48: CLI mtrace from-source Command Options (*continued*)

Option	Description
extra-hops <i>number</i>	(Optional) Sets the number of extra hops to trace past nonresponsive devices. Specify a value from 0 through 255 .
group <i>address</i>	(Optional) Traces the path for the specified group address. The default value is 192.0.2.0 .
interval <i>seconds</i>	(Optional) Sets the interval between statistics gathering. The default value is 10 .
max-hops <i>number</i>	(Optional) Sets the maximum number of hops to trace toward the source. Specify a value from 0 through 255 . The default value is 32 .
max-queries <i>number</i>	(Optional) Sets the maximum number of query attempts for any hop. Specify a value from 1 through 32 . The default value is 3 .
response <i>host</i>	(Optional) Sends the response packets to the specified hostname or IP address. By default, the response packets are sent to the device.
routing-instance <i>routing-instance-name</i>	(Optional) Traces the routing instance you specify.
ttl <i>number</i>	(Optional) Sets the time-to-live (TTL) value in the IP header of the query packets. Specify a hop count from 0 through 255 . The default value for local queries to the all routers multicast group is 1 . Otherwise, the default value is 127 .
wait-time <i>seconds</i>	(Optional) Sets the time to wait for a response packet. The default value is 3 seconds.
loop	(Optional) Loops indefinitely, displaying rate and loss statistics. To quit the mtrace command, press Ctrl-C.
multicast-response	(Optional) Forces the responses to use multicast.
unicast-response	(Optional) Forces the response packets to use unicast.
no-resolve	(Optional) Does not display hostnames.
no-router-alert	(Optional) Does not use the device alert IP option in the IP header.
brief	(Optional) Does not display packet rates and losses.
detail	(Optional) Displays packet rates and losses if a group address is specified.

The following is sample output from the **mtrace from-source** command:

```
user@host> mtrace from-source source 192.1.4.1 group 224.1.1.1
```

```
Mtrace from 192.1.4.1 to 192.1.30.2 via group 224.1.1.1 Querying full reverse
path... * * 0 ? (192.1.30.2) -1 ? (192.1.30.1) PIM thresh^ 1 -2
routerC.mycompany.net (192.1.40.2) PIM thresh^ 1 -3 hostA.mycompany.net
(192.1.4.1) Round trip time 22 ms; total ttl of 2 required. Waiting to accumulate
statistics...Results after 10 seconds: Source Response Dest Overall
```

```

      Packet Statistics For Traffic From 192.1.4.1 192.1.30.2
192.1.4.1 To 224.1.1.1      v      ___/  rtt  16 ms      Rate
Pct  Rate 192.168.195.37  192.1.40.2      routerC.mycompany.net
ttl    2              0/0  =  --    0 pps 192.1.40.1      192.1.30.1
?              v      \___  ttl    3      ?/0
0 pps 192.1.30.2      192.1.30.2  Receiver      Query Source

```

Each line of the trace display is usually in the following format (depending on the options selected and the responses from the devices along the path):

hop-number host (ip-address) protocolttl

Table 49 on page 393 summarizes the output fields of the display.



NOTE: The packet statistics gathered from Juniper Networks devices and routing nodes always display as 0.

Table 49: CLI mtrace from-source Command Output Summary

Field	Description
<i>hop-number</i>	Number of the hop (device) along the path.
<i>host</i>	Hostname, if available, or IP address of the device. If the no-resolve option was entered in the command, the hostname is not displayed.
<i>ip-address</i>	IP address of the device.
<i>protocol</i>	Protocol used.
<i>ttl</i>	TTL threshold.
Round trip time <i>milliseconds ms</i>	Total time between the sending of the query packet and the receiving of the response packet.
total ttl of <i>number</i> required	Total number of hops required to reach the source.
Source	Source IP address of the response packet.
Response Dest	Response destination IP address.
Overall	Average packet rate for all traffic at each hop.
Packet Statistics For Traffic From	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.
Receiver	IP address receiving the multicast packets.
Query Source	IP address of the host sending the query packets.

Related Documentation • [Monitoring Overview on page 7](#)

Monitoring Application Layer Gateways Features

- [Monitoring H.323 ALG Information on page 395](#)
- [Monitoring MGCP ALGs on page 396](#)
- [Monitoring SCCP ALGs on page 399](#)
- [Monitoring SIP ALGs on page 402](#)
- [Monitoring Voice ALG H.323 on page 406](#)
- [Monitoring Voice ALG MGCP on page 408](#)
- [Monitoring Voice ALG SCCP on page 411](#)
- [Monitoring Voice ALG SIP on page 414](#)
- [Monitoring Voice ALG Summary on page 419](#)

Monitoring H.323 ALG Information

Supported Platforms [SRX Series](#)

Purpose View the H.323 ALG counters information.

Action Select **Monitor>ALGs>H323** in the J-Web user interface, or enter the **show security alg h323 counters** command.

[Table 50 on page 395](#) summarizes key output fields in the H.323 counters display.

Table 50: Summary of Key H.323 Counters Output Fields

Field	Values	Additional Information
H.323 Counters Information		
Packets received	Number of H.323 ALG packets received.	—
Packets dropped	Number of H.323 ALG packets dropped.	—

Table 50: Summary of Key H.323 Counters Output Fields (*continued*)

Field	Values	Additional Information
RAS message received	Number of incoming RAS (Endpoint Registration, Admission, and Status) messages per second per gatekeeper received and processed.	—
Q.931 message received	Counter for Q.931 message received.	—
H.245 message received	Counter for H.245 message received.	—
Number of calls	Total number of H.323 ALG calls.	—
Number of active calls	Number of active H.323 ALG calls.	This counter displays the number of call legs and might not display the exact number of voice calls that are active. For instance, for a single active voice call between two endpoints, this counter might display a value of 2.
H.323 Error Counters		
Decoding errors	Number of decoding errors.	—
Message flood dropped	Error counter for message flood dropped.	—
NAT errors	H.323 ALG Network Address Translation (NAT) errors.	—
Resource manager errors	H.323 ALG resource manager errors.	—

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

Monitoring MGCP ALGs

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

This section contains the following topics:

- [Monitoring MGCP ALG Calls on page 397](#)
- [Monitoring MGCP ALG Counters on page 397](#)
- [Monitoring MGCP ALG Endpoints on page 399](#)

Monitoring MGCP ALG Calls

Purpose View information about MGCP ALG calls.

Action Select **Monitor>ALGs>MGCP>Calls** in the J-Web user interface. To view detailed information, select the endpoint on the MGCP calls page.

Alternatively, enter the **show security alg mgcp calls** command.

[Table 51 on page 397](#) summarizes key output fields in the MGCP calls display.

Table 51: Summary of Key MGCP Calls Output Fields

Field	Values	Additional Information
MGCP Calls Information		
Endpoint@GW	Endpoint name.	—
Zone	<ul style="list-style-type: none"> • trust—Trust zone. • untrust—Untrust zone. 	—
Call ID	Call identifier for ALG MGCP.	—
RM Group	Resource manager group ID.	—
Call Duration	Duration for which connection is active.	—
Connection Id	Connection identifier for MGCP ALG calls.	—
Calls Details: Endpoint		
Local SDP	IP address of the MGCP ALG local call owner, as per the Session Description Protocol (SDP).	—
Remote SDP	Remote IP address of the MGCP ALG remote call owner, as per the Session Description Protocol (SDP).	—

Monitoring MGCP ALG Counters

Purpose View MGCP ALG counters information.

Action Select **Monitor>ALGs>MGCP>Counters** in the J-Web user interface, or enter the **show security alg mgcp counters** command.

Table 52 on page 398 summarizes key output fields in the MGCP counters display.

Table 52: Summary of Key MGCP Counters Output Fields

Field	Values	Additional Information
MGCP Counters Information		
Packets received	Number of MGCP ALG packets received.	—
Packets dropped	Number of MGCP ALG packets dropped.	—
Message received	Number of MGCP ALG messages received.	—
Number of connections	Number of MGCP ALG connections.	—
Number of active connections	Number of active MGCP ALG connections.	—
Number of calls	Number of MGCP ALG calls.	—
Number of active calls	Number of MGCP ALG active calls.	—
Number of active transactions	Number of active transactions.	—
Number of re-transmission	Number of MGCP ALG retransmissions.	—
Error Counters		
Unknown-method	MGCP ALG unknown method errors.	—
Decoding error	MGCP ALG decoding errors.	—
Transaction error	MGCP ALG transaction errors.	—
Call error	MGCP ALG counter errors.	—
Connection error	MGCP ALG connection errors.	—
Connection flood drop	MGCP ALG connection flood drop errors.	—
Message flood drop	MGCP ALG message flood drop errors.	—
IP resolve error	MGCP ALG IP address resolution errors.	—
NAT error	MGCP ALG Network Address Translation (NAT) errors.	—
Resource manager error	MGCP ALG resource manager errors.	—

Monitoring MGCP ALG Endpoints

Purpose View information about MGCP ALG endpoints.

Action Select **Monitor>ALGs>MGCP>Endpoints** in the J-Web user interface. To view detailed information, select the gateway on the MGCP endpoints page.

Alternatively, enter the **show security alg mgcp endpoints** command.

Table 53 on page 399 summarizes key output fields in the MGCP endpoints display.

Table 53: Summary of Key MGCP Endpoints Output Fields

Field	Values	Additional Information
MGCP Endpoints		
Gateway	IP address of the gateway.	—
Zone	<ul style="list-style-type: none"> trust—Trust zone. untrust—Untrust zone. 	—
IP	IP address.	—
Endpoints: Gateway name		
Endpoint	Endpoint name.	—
Transaction #	Transaction identifier.	—
Call #	Call identifier.	—
Notified Entity	The certificate authority (CA) currently controlling the gateway.	—

Related Documentation

- [Monitoring Overview on page 7](#)
- [Monitoring Interfaces on page 426](#)

Monitoring SCCP ALGs

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

This section contains the following topics:

- [Monitoring SCCP ALG Calls on page 400](#)
- [Monitoring SCCP ALG Counters on page 400](#)

Monitoring SCCP ALG Calls

Purpose View information about SCCP ALG calls.

Action Select **Monitor>ALGs>SCCP>Calls** in the J-Web user interface. To view detailed information, select the client IP address on the SCCP calls page.

Alternatively, enter the **show security alg sccp calls** command.

[Table 54 on page 400](#) summarizes key output fields in the SCCP calls display.

Table 54: Summary of Key SCCP Calls Output Fields

Field	Values	Additional Information
SCCP Calls Information		
Client IP	IP address of the client.	—
Zone	Client zone identifier.	—
Call Manager	IP address of the call manager.	—
Conference ID	Conference call identifier.	—
RM Group	Resource manager group identifier.	—

Monitoring SCCP ALG Counters

Purpose View SCCP ALG counters information.

Action Select **Monitor>ALGs>SCCP>Count** in the J-Web user interface, or enter the **show security alg sccp counters** command.

[Table 55 on page 400](#) summarizes key output fields in the SCCP counters display.

Table 55: Summary of Key SCCP Counters Output Fields

Field	Values	Additional Information
SCCP Counters Information		
Clients currently registered	Number of SCCP ALG clients currently registered.	—
Active calls	Number of active SCCP ALG calls.	—
Total calls	Total number of SCCP ALG calls.	—

Table 55: Summary of Key SCCP Counters Output Fields (*continued*)

Field	Values	Additional Information
Packets received	Number of SCCP ALG packets received.	—
PDUs processed	Number of SCCP ALG protocol data units (PDUs) processed.	—
Current call rate	Number of calls per second.	—
Error counters		
Packets dropped	Number of packets dropped by the SCCP ALG.	—
Decode errors	SCCP ALG decoding errors.	—
Protocol errors	Number of protocol errors.	—
Address translation errors	Number of Network Address Translation (NAT) errors encountered by SCCP ALG.	—
Policy lookup errors	Number of packets dropped because of a failed policy lookup.	—
Unknown PDUs	Number of unknown protocol data units (PDUs).	—
Maximum calls exceed	Number of times the maximum SCCP calls limit was exceeded.	—
Maximum call rate exceed	Number of times the maximum SCCP call rate exceeded.	—
Initialization errors	Number of initialization errors.	—
Internal errors	Number of internal errors.	—
Unsupported feature	Number of unsupported feature errors.	—
Non specific error	Number of nonspecific errors.	—

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

Monitoring SIP ALGs

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring SIP ALG Calls on page 402](#)
- [Monitoring SIP ALG Counters on page 402](#)
- [Monitoring SIP ALG Rate Information on page 404](#)
- [Monitoring SIP ALG Transactions on page 405](#)

Monitoring SIP ALG Calls

Purpose View information about SIP ALG calls.

Action Select **Monitor>ALGs>SIP>Calls** in the J-Web user interface. To view detailed information, select the Call Leg on the SIP calls page.

Alternatively, enter the **show security alg sip calls detail** command.

[Table 56 on page 402](#) summarizes key output fields in the SIP calls display.

Table 56: Summary of Key SIP Calls Output Fields

Field	Values	Additional Information
SIP Calls Information		
Call Leg	Call length identifier.	—
Zone	Client zone identifier.	—
RM Group	Resource manager group identifier.	—
Local Tag	Local tag for the SIP ALG User Agent server.	—
Remote Tag	Remote tag for the SIP ALG User Agent server.	—

Monitoring SIP ALG Counters

Purpose View SIP ALG counters information.

Action Select **Monitor>ALGs>SIP>Count** in the J-Web user interface, or enter the **show security alg sip counters** command.

Table 57 on page 403 summarizes key output fields in the SIP counters display.

Table 57: Summary of Key SIP Counters Output Fields

Field	Values	Additional Information
SIP Counters Information		
INVITE	Number of INVITE requests sent.	An INVITE request is sent to invite another user to participate in a session.
CANCEL	Number of CANCEL requests sent.	A user can send a CANCEL request to cancel a pending INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE had sent a final response for the INVITE before it received the CANCEL.
ACK	Number of ACK requests sent.	The user from whom the INVITE originated sends an ACK request to confirm reception of the final response to the INVITE request.
BYE	Number of BYE requests sent.	A user sends a BYE request to abandon a session. A BYE request from either user automatically terminates the session.
REGISTER	Number of REGISTER requests sent.	A user sends a REGISTER request to a SIP registrar server to inform it of the current location of the user. A SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server attempting to locate a user.
OPTIONS	Number of OPTIONS requests sent.	An OPTION message is used by the User Agent (UA) to obtain information about the capabilities of the SIP proxy. A server responds with information about what methods, session description protocols, and message encoding it supports.
INFO	Number of INFO requests sent.	An INFO message is used to communicate mid-session signaling information along the signaling path for the call.
MESSAGE	Number of MESSAGE requests sent.	SIP messages consist of requests from a client to a server and responses to the requests from a server to a client with the purpose of establishing a session (or a call).
NOTIFY	Number of NOTIFY requests sent.	A NOTIFY message is sent to inform subscribers of changes in state to which the subscriber has a subscription.
REFER	Number of REFER requests sent.	A REFER request is used to refer the recipient (identified by the Request-URI) to a third party by the contact information provided in the request.
SUBSCRIBE	Number of SUBSCRIBE requests sent.	A SUBSCRIBE request is used to request current state and state updates from a remote node.

Table 57: Summary of Key SIP Counters Output Fields (*continued*)

Field	Values	Additional Information
UPDATE	Number of UPDATE requests sent.	An UPDATE request is used to create a temporary opening in the firewall (pinhole) for new or updated Session Description Protocol (SDP) information. The following header fields are modified: Via, From, To, Call-ID, Contact, Route, and Record-Route.
SIP Error Counters		
Total Pkt-in	SIP ALG total packets received.	—
Total Pkt dropped on error	Number of packets dropped by the SIP ALG.	—
Transaction error	SIP ALG transaction errors.	—
Call error	SIP ALG call errors.	—
IP resolve error	SIP ALG IP address resolution errors.	—
NAT error	SIP ALG NAT errors.	—
Resource manager error	SIP ALG resource manager errors.	—
RR header exceeded max	Number of times the SIP ALG RR (Record-Route) headers exceeded the maximum limit.	—
Contact header exceeded max	Number of times the SIP ALG contact header exceeded the maximum limit.	—
Call dropped due to limit	SIP ALG calls dropped because of call limits.	—
SIP stack error	SIP ALG stack errors.	—

Monitoring SIP ALG Rate Information

Purpose View SIP ALG rate information.

Action Select **Monitor>ALGs>SIP>Rate** in the J-Web user interface, or enter the **show security alg sip rate** command.

Table 58 on page 405 summarizes key output fields in the SIP rate display.

Table 58: Summary of Key SIP Rate Output Fields

Field	Values	Additional Information
SIP Rate Information		
CPU ticks per microseconds is	SIP ALG CPU ticks per microsecond.	—
Time taken for the last message in microseconds is	Time, in microseconds, that the last SIP ALG message needed to transit the network.	—
Number of messages in 10 minutes	Total number of SIP ALG messages transiting the network in 10 minutes.	—
Time taken by the messages in 10 minutes	Total time, in microseconds, during an interval of less than 10 minutes for the specified number of SIP ALG messages to transit the network.	—
Rate	Number of SIP ALG messages per second transiting the network.	—

Monitoring SIP ALG Transactions

Purpose View information about SIP ALG transactions.

Action Select **Monitor>ALGs>SIP>Transactions** in the J-Web user interface, or enter the **show security alg sip transactions** command.

Table 59 on page 405 summarizes key output fields in the SIP transactions display.

Table 59: Summary of Key SIP Transactions Output Fields

Field	Values	Additional Information
SIP Transactions Information		
Transaction Name	<ul style="list-style-type: none"> • UAS—SIP ALG User Agent server transaction name. • UAC—SIP ALG User Agent client transaction name. 	—

Table 59: Summary of Key SIP Transactions Output Fields (*continued*)

Field	Values	Additional Information
Method	<p>The method to be performed on the resource. Possible methods:</p> <ul style="list-style-type: none"> • INVITE—Initiate call • ACK—Confirm final response • BYE—Terminate and transfer call • CANCEL—Cancel searches and “ringing” • OPTIONS—Features support by the other side • REGISTER—Register with location service 	—

Related Documentation

- [Monitoring Overview on page 7](#)
- [Monitoring Interfaces on page 426](#)

Monitoring Voice ALG H.323

Supported Platforms [SRX Series](#)

Purpose Use the monitoring functionality to view the ALG H.323 page.

Action To monitor ALG H.323 select **Monitor>Security>Voice ALGs>H.323** in the J-Web user interface.

Meaning [Table 60 on page 406](#) summarizes key output fields in the ALG H.323 page.

Table 60: ALG H.323 Monitoring Page

Field	Value	Additional Information
Virtual Chassis Member	Display the list of virtual chassis member.	Select one of the virtual chassis members listed.
Refresh Interval (30 sec)	Displays the time interval set for page refresh.	Select the time interval from the drop-down list.
Refresh	Displays the option to refresh the page.	—
Clear	Provides an option to clear the monitor summary.	Click clear to clear the monitor summary.

H.323 Counter Summary

Table 60: ALG H.323 Monitoring Page (*continued*)

Field	Value	Additional Information
Category	Displays the following categories: <ul style="list-style-type: none"> • Packets received—Number of ALG H.323 packets received. • Packets dropped—Number of ALG H.323 packets dropped. • RAS message received—Number of incoming RAS (Registration, Admission, and Status) messages per second per gatekeeper received and processed. • Q.931 message received—Counter for Q.931 message received. • H.245 message received—Counter for H.245 message received. • Number of calls—Total number of ALG H.323 calls. • Number of active calls—Number of active ALG H.323 calls. • Number of DSCP Marked—Number of DSCP Marked on ALG H.323 calls. 	—
Count	Provides count of response codes for each H.323 counter summary category.	—
H.323 Error Counter		
Category	Displays the following categories: <ul style="list-style-type: none"> • Decoding errors—Number of decoding errors. • Message flood dropped—Error counter for message flood dropped. • NAT errors—H.323 ALG NAT errors. • Resource manager errors—H.323 ALG resource manager errors. • DSCP Marked errors—H.323 ALG DSCP marked errors. 	—
Count	Provides count of response codes for each H.323 error counter category.	—
Counter Summary Chart		
Packets Received	Provides the graphical representation of the packets received.	—
H.323 Message Counter		

Table 60: ALG H.323 Monitoring Page (*continued*)

Field	Value	Additional Information
Category	Displays the following categories: <ul style="list-style-type: none"> • RRQ—Registration Request message counter. • RCF—Registration Confirmation Message. • ARQ—Admission Request message counter. • ACF—Admission Confirmation • URQ—Unregistration Request. • UCF—Unregistration Confirmation. • DRQ—Disengage Request. • DCF—Disengage Confirmation. • Oth RAS—Other incoming Registration, Admission, and Status messages message counter. • Setup—Timeout value, in seconds, for the response of the outgoing setup message. • Alert—Alert message type. • Connect—Connect setup process. • CallProd—Number of call production messages sent. • Info—Number of info requests sent. • RelCmpl—Number of Rel Cmpl message ssent. • Facility—Number of facility messages sent. • Empty—Empty capabilities to the support message counter. • OLC—Open Local Channel message counter. • OLC ACK—Open Local Channel Acknowledge message counter. • Oth H245—Other H.245 message counter 	—
Count	Provides count of response codes for each H.323 message counter category.	—

- Related Documentation**
- [Monitoring Voice ALG Summary on page 419](#)
 - [Monitoring Voice ALG MGCP on page 408](#)
 - [Monitoring Voice ALG SCCP on page 411](#)
 - [Monitoring Voice ALG SIP on page 414](#)

Monitoring Voice ALG MGCP

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [vSRX](#)

Purpose Use the monitoring functionality to view the voice ALG MGCP page.

Action To monitor ALG MGCP, select **Monitor>Security>Voice ALGs>MGCP** in the J-Web user interface.

Meaning Table 61 on page 409 summarizes key output fields in the voice ALG MGCP page.

Table 61: Voice ALG MGCP Monitoring Page

Field	Value	Additional Information
Virtual Chassis Member	Displays the list of virtual chassis member.	Select one of the virtual chassis members listed.
Refresh Interval (30 sec)	Displays the time interval set for page refresh.	Select the time interval from the drop-down list.
Refresh	Displays the option to refresh the page.	—
Clear	Provides an option to clear the monitor summary.	Click Clear to clear the monitor summary.
Counters		
MGCP Counters Summary		
Category	Displays the following categories: <ul style="list-style-type: none"> • Packets Received—Number of ALG MGCP packets received. • Packets Dropped— Number of ALG MGCP packets dropped. • Message received— Number of ALG MGCP messages received. • Number of connections— Number of ALG MGCP connections. • Number of active connections— Number of active ALG MGCP connections. • Number of calls— Number of ALG MGCP calls. • Number of active calls— Number of active ALG MGCP calls. • Number of active transactions— Number of active transactions. • Number of transactions— Number of transactions. • Number of re-transmission—Number of ALG MGCP retransmissions. • Number of active endpoints— Number of MGCP active endpoints. • Number of DSCP marked— Number of MGCP DSCPs marked. 	—
Count	Provides the count of response codes for each MGCP counter summary category.	—

Table 61: Voice ALG MGCP Monitoring Page (*continued*)

Field	Value	Additional Information
MGCP Error Counter		
Category	Displays the following categories: <ul style="list-style-type: none"> • Unknown-method— MGCP ALG unknown method errors. • Decoding error— MGCP ALG decoding errors. • Transaction error— MGCP ALG transaction errors. • Call error— MGCP ALG call counter errors. • Connection error— MGCP ALG connection errors. • Connection flood drop— MGCP ALG connection flood drop errors. • Message flood drop— MGCP ALG message flood drop error. • IP resolve error— MGCP ALG IP address resolution errors. • NAT error— MGCP ALG NAT errors. • Resource manager error— MGCP ALG resource manager errors. • DSCP Marked error— MGCP ALG DSCP marked errors. 	—
Count	Provides the count of response codes for each summary error counter category.	—
Counter Summary Chart	Displays the Counter Summary Chart.	—
MGCP Packet Counters		
Category	Displays the following categories: <ul style="list-style-type: none"> • CRCX— Create Connection • MDCX— Modify Connection • DLCX— Delete Connection • AUEP— Audit Endpoint • AUCX— Audit Connection • NTFY— Notify MGCP • RSIP— Restart in Progress • EPCF— Endpoint Configuration • RQNT— Request for Notification • 000-199— Respond code is 0-199 • 200-299— Respond code is 200-299 • 300-399— Respond code is 300-399 	—
Count	Provides count of response codes for each MGCP packet counter category.	—

Table 61: Voice ALG MGCP Monitoring Page (*continued*)

Field	Value	Additional Information
Calls		
Endpoint@GW	Displays the endpoint name.	—
Zone	Displays the following options: <ul style="list-style-type: none"> • trust—Trust zone. • untrust—Untrust zone. 	—
Endpoint IP	Displays the endpoint IP address.	—
Call ID	Displays the call identifier for ALG MGCP.	—
RM Group	Displays the resource manager group ID.	—
Call Duration	Displays the duration for which connection is active.	—

- Related Documentation**
- [Monitoring Voice ALG Summary on page 419](#)
 - [Monitoring Voice ALG H.323 on page 406](#)
 - [Monitoring Voice ALG SCCP on page 411](#)
 - [Monitoring Voice ALG SIP on page 414](#)

Monitoring Voice ALG SCCP

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

Purpose Use the monitoring functionality to view the voice ALG SCCP page.

Action To monitor voice ALG SCCP, select **Monitor>Security>Voice ALGs>SCCP** in the J-Web user interface.

Meaning [Table 62 on page 411](#) summarizes key output fields in the voice ALG SCCP page.

Table 62: Voice ALG SCCP Monitoring Page

Field	Value	Additional Information
Virtual Chassis Member	Displays the list of virtual chassis member.	Select one of the virtual chassis members listed.
Refresh Interval (30 sec)	Displays the time interval set for page refresh.	Select the time interval from the drop-down list.

Table 62: Voice ALG SCCP Monitoring Page (*continued*)

Field	Value	Additional Information
Refresh	Displays the option to refresh the page.	—
Clear	Provides an option to clear the monitor summary.	Click Clear to clear the monitor summary.
SCCP Call Statistics		
Category	Displays the following categories: <ul style="list-style-type: none"> • Active client sessions— Number of active SCCP ALG client sessions. • Active calls— Number of active SCCP ALG calls. • Total calls— Total number of SCCP ALG calls. • Packets received— Number of SCCP ALG packets received. • PDUs processed— Number of SCCP ALG protocol data units (PDUs) processed. • Current call rate— Number of calls per second. • DSCPs Marked— Number of DSCP marked. 	—
Count	Provides count of response codes for each SCCP call statistics category.	—
Call Statistics Chart	Displays the Call Statistics chart.	—
SCCP Error Counters		

Table 62: Voice ALG SCCP Monitoring Page (*continued*)

Field	Value	Additional Information
Category	Displays the following categories: <ul style="list-style-type: none"> • Packets dropped— Number of packets dropped by the SCCP ALG. • Decode errors— Number of SCCP ALG decoding errors. • Protocol errors— Number of protocol errors. • Address translation errors— Number of NAT errors encountered by SCCP ALG. • Policy lookup errors— Number of packets dropped because of a failed policy lookup. • Unknown PDUs— Number of unknown PDUs. • Maximum calls exceed— Number of times the maximum SCCP calls limit was exceeded. • Maximum call rate exceed— Number of times the maximum SCCP call rate was exceeded. • Initialization errors— Number of initialization errors. • Internal errors— Number of internal errors. • Nonspecific errors— Number of nonspecific errors. • No active calls to be deleted— Number of no active calls to be deleted. • No active client sessions to be deleted— Number of no active client sessions to be deleted. • Session cookie created error— Number of session cookie created errors. • Invalid NAT cookies deleted— Number of invalid NAT cookies deleted. • NAT cookies not found— Number of NAT cookies not found. • DSCP Marked Error— Number of DSCP marked errors. 	—
Count	Provides count of response codes for each SCCP error counter category.	—
Calls		
Client IP	Displays the IP address of the client.	—
Zone	Displays the client zone identifier.	—
Call Manager	Displays the IP address of the call manager.	—
Conference ID	Displays the conference call identifier.	—
RM Group	Displays the resource manager group identifier.	—

- Related Documentation**
- [Monitoring Voice ALG Summary on page 419](#)
 - [Monitoring Voice ALG H.323 on page 406](#)
 - [Monitoring Voice ALG MGCP on page 408](#)
 - [Monitoring Voice ALG SIP on page 414](#)

Monitoring Voice ALG SIP

Supported Platforms [SRX Series, vSRX](#)

Purpose Use the monitoring functionality to view the voice ALG SIP page.

Action To monitor voice ALG SIP select **Monitor>Security>Voice ALGs>SIP** in the J-Web user interface.

Meaning [Table 63 on page 414](#) summarizes key output fields in the voice ALG SIP page.

Table 63: Voice ALG SIP Monitoring Page

Field	Value	Additional Information
Virtual Chassis Member	Displays the list of virtual chassis members.	Select one of the virtual chassis members listed.
Refresh Interval (30 sec)	Displays the time interval set for page refresh.	Select the time interval from the drop-down list.
Refresh	Displays the option to refresh the page.	—
Clear	Provides an option to clear the monitor summary.	Click Clear to clear the monitor summary.

Counters

SIP Counters Information

Table 63: Voice ALG SIP Monitoring Page (*continued*)

Field	Value	Additional Information
Method	<p>Displays the SIP counter information. The available options are:</p> <ul style="list-style-type: none"> • BYE— Number of BYE requests sent. A user sends a BYE request to abandon a session. A BYE request from either user automatically terminates the session. • REGISTER— Number of REGISTER requests sent. A user sends a REGISTER request to a SIP registrar server to inform it of the current location of the user. The SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server attempting to locate a user. • OPTIONS— Number of OPTIONS requests sent. An OPTION message is used by the User Agent (UA) to obtain information about the capabilities of the SIP proxy. A server responds with information about what methods, session description protocols, and message encoding it supports. • INFO— Number of INFO requests sent. An INFO message is used to communicate mid-session signaling information along the signaling path for the call. • MESSAGE— Number of MESSAGE requests sent. SIP messages consist of requests from a client to the server and responses to the requests from the server to a client for the purpose of establishing a session (or a call). 	—

SIP Counters Information (*continued*)

Table 63: Voice ALG SIP Monitoring Page (*continued*)

Field	Value	Additional Information
Method	<ul style="list-style-type: none"> • NOTIFY— Number of NOTIFY requests sent. A NOTIFY message is sent to inform subscribers about the change in state of the subscription. • PRACK— Number of PRACK requests sent. The PRACK request plays the same role as the ACK request, but for provisional responses. • PUBLISH— Number of PUBLISH requests sent. The PUBLISH request is used for publishing the event state. PUBLISH is similar to REGISTER that allows a user to create, modify, and remove state in another entity which manages this state on behalf of the user. • REFER— Number of REFER requests sent. A REFER request is used to refer the recipient (identified by the Request-URI) to a third party identified by the contact information provided in the request. • SUBSCRIBE— Number of SUBSCRIBE requests sent. A SUBSCRIBE request is used to request current state and state information updates from a remote node. • UPDATE— Number of UPDATE requests sent. An UPDATE request is used to create a temporary opening in the firewall (pinhole) for new or updated Session Description Protocol (SDP) information. The following header fields are modified: Via, From, To, Call-ID, Contact, Route, and Record-Route. • BENOTIFY— Number of BENOTIFY requests sent. A BENOTIFY request is used to reduce the unnecessary SIP signaling traffic on application servers. Applications that do not need a response for a NOTIFY request can enhance performance by enabling BENOTIFY. • SERVICE— Number of SERVICE requests sent. The SERVICE method is used by a SIP client to request a service from a SIP server. It is a standard SIP message and will be forwarded until it reaches the server or end user that is performing the service. • OTHER— Number of OTHER requests sent. 	—
T, RT	Displays the transmit and retransmit method.	—
1xx, RT	Displays one transmit and retransmit method.	—
2xx, RT	Displays two transmit and retransmit methods.	—
3xx, RT	Displays three transmit and retransmit methods.	—
4xx, RT	Displays four transmit and retransmit methods.	—
5xx, RT	Displays five transmit and retransmit methods.	—
6xx, RT	Displays six transmit and retransmit methods.	—
Calls		
Call ID	Displays the call ID.	—

Table 63: Voice ALG SIP Monitoring Page (*continued*)

Field	Value	Additional Information
Method	Displays the call method used.	—
State	Displays the state of the ALG SIP.	—
Group ID	Displays the group identifier.	—
Invite Method Chart	Displays the invite method chart. The available options are: <ul style="list-style-type: none"> • T/RT • 1xx/ RT • 2xx/ RT • 3xx/ RT • 4xx/ RT • 5xx/ RT • 6xx/ RT 	—

SIP Error Counters

Table 63: Voice ALG SIP Monitoring Page (*continued*)

Field	Value	Additional Information
Category	<p>Displays the SIP error counters. The available options are:</p> <ul style="list-style-type: none"> • Total Pkt-in— Number of SIP ALG total packets received. • Total Pkt dropped on error— Number of packets dropped by the SIP ALG. • Call error— SIP Number of ALG call errors. • IP resolve error— Number of SIP ALG IP address resolution errors. • NAT error— SIP Number of ALG NAT errors. • Resource manager error— Number of SIP ALG resource manager errors. • RR header exceeded max— Number of times the SIP ALG RR (Record-Route) headers exceeded the maximum limit. • Contact header exceeded max— Number of times the SIP ALG contact header exceeded the maximum limit. • Call dropped due to limit— Number of SIP ALG calls dropped because of call limits. • SIP stack error— Number of SIP ALG stack errors. • SIP Decode error— Number of SIP ALG decode errors. • SIP unknown method error— Number of SIP ALG unknow method errors. • SIP DSCP marked—SIP ALG DSCP marked. • SIP DSCP marked error— Number of SIP ALG DSCPs marked. • RTO message sent— Number of SIP ALG marked RTO messages sent. • RTO message received— Number of SIP ALG RTO messages received. • RTO buffer allocation failure— Number of SIP ALG RTO buffer allocation failures. • RTO buffer transmit failure— Number of SIP ALG RTO buffer transmit failures. • RTO send processing error— Number of SIP ALG RTO send processing errors. • RTO receiving processing error— Number of SIP ALG RTO receiving processing errors. • RTO receive invalid length— Number of SIP ALG RTOs receiving invalid length. • RTO receive call process error— Number of SIP ALG RTO receiving call process errors. • RTO receive call allocation error— Number of SIP ALG RTO receiving call allocation error. • RTO receive call register error— Number of SIP ALG RTO receiving call register errors. • RTO receive invalid status error— Number of SIP ALG RTO receiving register errors. 	—
Count	Provides count of response codes for each SIP ALG counter category.	—

- Related Documentation**
- [Monitoring Voice ALG Summary on page 419](#)
 - [Monitoring Voice ALG H.323 on page 406](#)
 - [Monitoring Voice ALG MGCP on page 408](#)
 - [Monitoring Voice ALG SCCP on page 411](#)

Monitoring Voice ALG Summary

Supported Platforms [SRX Series, vSRX](#)

Purpose Use the monitoring functionality to view the voice ALG summary page.

Action To monitor voice ALG summary, select **Monitor>Security>Voice ALGs>Summary** in the J-Web user interface.

Meaning [Table 64 on page 419](#) summarizes key output fields in the voice ALG summary page.

Table 64: Voice ALG Summary Monitoring Page

Field	Value	Additional Information
Virtual Chassis Member	Display the list of virtual chassis member.	Select one of the virtual chassis members listed.
Refresh Interval (30 sec)	Displays the time interval set for page refresh.	Select the time interval from the drop-down list.
Refresh	Displays the option to refresh the page.	—
Clear	Provides an option to clear the monitor summary.	Click Clear to clear the monitor summary.
Protocol Name	Displays the protocols configured.	—
Total Calls	Displays the total number of calls.	—
Number of Active Calls	Displays the number of active calls.	—
Number of Received Packets	Displays the number of packets received.	—
Number of Errors	Displays the number of errors.	—
H.323 Calls Chart	Displays the H.323 calls chart.	—
MGCP Calls Chart	Displays the MGCP calls chart.	—
SCCP Calls Chart	Displays the SCCP calls chart.	—
SIP Calls Chart	Displays the SIP calls chart.	—

**Related
Documentation**

- [Monitoring Voice ALG H.323 on page 406](#)
- [Monitoring Voice ALG MGCP on page 408](#)
- [Monitoring Voice ALG SCCP on page 411](#)
- [Monitoring Voice ALG SIP on page 414](#)

CHAPTER 24

Monitoring Interfaces and Switching Functions

- [Displaying Real-Time Interface Information on page 421](#)
- [Monitoring Address Pools on page 423](#)
- [Monitoring Ethernet Switching on page 424](#)
- [Monitoring GVRP on page 425](#)
- [Monitoring Interfaces on page 426](#)
- [Monitoring MPLS Traffic Engineering Information on page 427](#)
- [Monitoring PPP on page 433](#)
- [Monitoring PPPoE on page 434](#)
- [Monitoring Spanning Tree on page 438](#)
- [Monitoring the WAN Acceleration Interface on page 439](#)

Displaying Real-Time Interface Information

Supported Platforms [SRX Series, vSRX](#)

Enter the **monitor interface** command to display real-time traffic, error, alarm, and filter statistics about a physical or logical interface:

```
user@host> monitor interface (interface-name | traffic)
```

Replace ***interface-name*** with the name of a physical or logical interface. If you specify the **traffic** option, statistics for all active interfaces display.

The real-time statistics update every second. The **Current delta** and **Delta** columns display the amount the statistics counters have changed since the **monitor interface** command was entered or since you cleared the delta counters. [Table 65 on page 422](#) and [Table 66 on page 422](#) list the keys you use to control the display using the ***interface-name*** and **traffic** options. (The keys are not case sensitive.)

Table 65: CLI monitor interface Output Control Keys

Key	Action
c	Clears (returns to 0) the delta counters in the Current delta column. The statistics counters are not cleared.
f	Freezes the display, halting the update of the statistics and delta counters.
i	Displays information about a different interface. You are prompted for the name of a specific interface.
n	Displays information about the next interface. The device scrolls through the physical and logical interfaces in the same order in which they are displayed by the show interfaces terse command.
q or ESC	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

Table 66: CLI monitor interface traffic Output Control Keys

Key	Action
b	Displays the statistics in units of bytes and bytes per second (bps).
c	Clears (returns to 0) the delta counters in the Delta column. The statistics counters are not cleared.
d	Displays the Delta column instead of the rate column—in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or ESC	Quits the command and returns to the command prompt.
r	Displays the rate column—in bps and pps—instead of the Delta column.

The following are sample displays from the **monitor interface** command:

```
user@host> monitor interface fe-0/0/0
```

```

host1                               Seconds: 5                               Time: 04:38:40
                                      Delay: 3/0/10

Interface: fe-0/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 1000mbps
Traffic statistics:
  Input bytes:           885405423 (3248 bps)           [2631]
  Output bytes:          137411893 (3344 bps)           [10243]
  Input packets:         7155064 (2 pps)                [28]
  Output packets:        636071 (1 pps)                 [23]
Error statistics:
  Input errors:          0                               [0]
```



```

Input drops:                                0                [0]
Input framing errors:                       0                [0]
Policed discards:                           0                [0]
L3 incompletes:                             0                [0]
L2 channel errors:                           0                [0]
L2 mismatch timeouts:                       0                [0]
Carrier transitions:                         1                [0]
Output errors:                              0                [0]
Output drops:                               0                [0]
Aged packets:                               0                [0]
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets                           73083              [16]
  Broadcast packets                         3629058            [5]
  Multicast packets                         3511364            [3]
  Oversized frames                          0                  [0]
  Packet reject count                       0                  [0]
  DA rejects                               0                  [0]
  SA rejects                               0                  [0]
Output MAC/Filter Statistics:
  Unicast packets                           629555             [28]
  Broadcast packets                         6494 Multicast packet [0]

```



NOTE: The output fields that display when you enter the monitor interface *interface-name* command are determined by the interface you specify.

```
user@host> monitor interface traffic
```

Interface	Link	Input packets	(pps)	Output packets	(pps)
fe-0/0/0	Up	42334	(5)	23306	(3)
fe-0/0/1	Up	587525876	(12252)	589621478	(12891)

Related Documentation • [Monitoring Interfaces on page 426](#)

Monitoring Address Pools

Supported Platforms [SRX Series, vSRX](#)

Purpose Use the monitoring functionality to view the Address Pools page.

Action To monitor Address Pools, select **Monitor>Access>Address Pools** in the J-Web user interface.

Meaning [Table 67 on page 423](#) summarizes key output fields in the Address Pools page.

Table 67: Address Pools Monitoring Page

Field	Values	Additional Information
Address Pool Properties		

Table 67: Address Pools Monitoring Page (*continued*)

Field	Values	Additional Information
Address Pool Name	Displays the name of the address pool.	-
Network Address	Displays the IP network address of the address pool.	-
Address Ranges	Displays the name, the lower limit, and the upper limit of the address range.	-
Primary DNS	Displays the primary-dns IP address.	-
Secondary DNS	Displays the secondary-dns IP address.	-
Primary WINS	Displays the primary-wins IP address.	-
Secondary WINS	Displays the secondary-wins IP address.	-
Address Pool Address Assignment		
IP Address	Displays the IP address of the address pool.	-
Hardware Address	Displays the hardware MAC address of the address pool.	-
Host/User	Displays the user name using the address pool.	-
Type	Displays the authentication type used by the address pool	The authentication types can be extended authentication (XAuth) or IKE Authentication.

- Related Documentation**
- [Monitoring Interfaces on page 426](#)
 - [Threats Monitoring Report on page 494](#)

Monitoring Ethernet Switching

Supported Platforms [SRX Series, vSRX](#)

Purpose View information about the Ethernet Switching interface details.

Action Select **Monitor>Switching>Ethernet Switching** in the J-Web user interface, or enter the following CLI commands:

- **show ethernet-switching table**
- **show ethernet-switching mac-learning-log**

[Table 68 on page 425](#) summarizes the Ethernet Switching output fields.

Table 68: Summary of Ethernet Switching Output Fields

Field	Values	Additional Information
VLAN	The VLAN for which Ethernet Switching is enabled.	-
MAC Address	The MAC address associated with the VLAN. If a VLAN range has been configured for a VLAN, the output displays the MAC addresses for the entire series of VLANs that were created with that name.	-
Type	The type of MAC address. Values are: <ul style="list-style-type: none"> static—The MAC address is manually created. learn—The MAC address is learned dynamically from a packet's source MAC address. flood—The MAC address is unknown and flooded to all members. 	-
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	-
Interfaces	Interface associated with learned MAC addresses or All-members (flood entry).	-
VLAN-ID	The VLAN ID.	-
MAC Address	The learned MAC address.	-
Time	Timestamp when the MAC address was added or deleted from the log.	-
State	Indicates the MAC address learned on the interface.	-

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

Monitoring GVRP

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [vSRX](#)

Purpose Use the monitoring functionality to view the GVRP page.

Action To monitor GVRP select **Monitor>Switching>GVRP** in the J-Web user interface.

Meaning [Table 69 on page 426](#) summarizes key output fields in the GVRP page.

Table 69: GVRP Monitoring Page

Field	Value	Additional Information
Global GVRP Configuration		
GVRP Status	Displays whether GVRP is enabled or disabled.	—
GVRP Timer	Displays the GVRP timer in millisecond.	—
Join	The number of milliseconds the interfaces must wait before sending VLAN advertisements.	—
Leave	The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message.	—
Leave All	The interval in milliseconds at which Leave All messages are sent on interfaces. Leave All messages maintain current GVRP VLAN membership information in the network.	—
GVRP Interface Details		
Interface Name	The interface on which GVRP is configured.	—
Protocol Status	Displays whether GVRP is enabled or disabled.	—

- Related Documentation**
- [Monitoring Ethernet Switching on page 424](#)
 - [Monitoring Spanning Tree on page 438](#)

Monitoring Interfaces

Supported Platforms [SRX Series, vSRX](#)

Purpose View general information about all physical and logical interfaces for a device.

Action Select **Monitor>Interfaces** in the J-Web user interface . The J-Web Interfaces page displays the following details about each device interface:

- Port—Indicates the interface name.
- Admin Status—Indicates whether the interface is enabled (Up) or disabled (Down).
- Link Status—Indicates whether the interface is linked (Up) or not linked (Down).
- Address—Indicates the IP address of the interface.
- Zone—Indicates whether the zone is an untrust zone or a trust zone.

- Services—Indicates services that are enabled on the device, such as HTTP and SSH.
- Protocols—Indicates protocols that are enabled on the device, such as BGP and IGMP.
- Input Rate graph—Displays interface bandwidth utilization. Input rates are shown in bytes per second.
- Output Rate graph—Displays interface bandwidth utilization. Output rates are shown in bytes per second.
- Error Counters chart—Displays input and output error counters in the form of a bar chart.
- Packet Counters chart—Displays the number of broadcast, unicast, and multicast packet counters in the form of a pie chart. (Packet counter charts are supported only for interfaces that support MAC statistics.)

To change the interface display, use the following options:

- Port for FPC—Controls the member for which information is displayed.
- Start/Stop button—Starts or stops monitoring the selected interfaces.
- Show Graph—Displays input and output packet counters and error counters in the form of charts.
- Pop-up button—Displays the interface graphs in a separate pop-up window.
- Details—Displays extensive statistics about the selected interface, including its general status, traffic information, IP address, I/O errors, class-of-service data, and statistics.
- Refresh Interval—Indicates the duration of time after which you want the data on the page to be refreshed.
- Clear Statistics—Clears the statistics for the selected interface.

Alternatively, you can enter the following **show** commands in the CLI to view interface status and traffic statistics:

- **show interfaces terse**



NOTE: On SRX Series devices, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

- **show interfaces detail**
- **show interfaces extensive**
- **show interfaces *interface-name***

Monitoring MPLS Traffic Engineering Information

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

This section contains the following topics:

- [Monitoring MPLS Interfaces on page 428](#)
- [Monitoring MPLS LSP Information on page 428](#)
- [Monitoring MPLS LSP Statistics on page 430](#)
- [Monitoring RSVP Session Information on page 431](#)
- [Monitoring MPLS RSVP Interfaces Information on page 432](#)

Monitoring MPLS Interfaces

Purpose View the interfaces on which MPLS is configured, including operational state and any administrative groups applied to an interface.

Action Select **Monitor>MPLS>Interfaces** in the J-Web user interface, or enter the **show mpls interface** command.

[Table 70 on page 428](#) summarizes key output fields in the MPLS interface information display.

Table 70: Summary of Key MPLS Interface Information Output Fields

Field	Values	Additional Information
Interface	Name of the interface on which MPLS is configured.	—
State	State of the specified interface: Up or Dn (down).	—
Administrative groups	Administratively assigned colors of the MPLS link configured on the interface.	—

See Also

- [Monitoring MPLS LSP Information on page 428](#)
- [Monitoring MPLS LSP Statistics on page 430](#)
- [Monitoring RSVP Session Information on page 431](#)
- [Monitoring MPLS RSVP Interfaces Information on page 432](#)

Monitoring MPLS LSP Information

Purpose View all label-switched paths (LSPs) configured on the services router, including all inbound (ingress), outbound (egress), and transit LSP information.

Action Select **Monitor>MPLS>LSP Information** in the J-Web user interface, or enter the **show mpls lsp** command.

[Table 71 on page 429](#) summarizes key output fields in the MPLS LSP information display.

Table 71: Summary of Key MPLS LSP Information Output Fields

Field	Values	Additional Information
Ingress LSP	Information about LSPs on the inbound device. Each session has one line of output.	–
Egress LSP	Information about the LSPs on the outbound device. Each session has one line of output.	MPLS learns this information by querying RSVP, which holds all the transit and outbound session information.
Transit LSP	Number of LSPs on the transit routers and the state of these paths.	MPLS learns this information by querying RSVP, which holds all the transit and outbound session information.
To	Destination (outbound device) of the session.	–
From	Source (inbound device) of the session.	–
State	State of the path. It can be Up , Down , or AdminDn .	AdminDn indicates that the LSP is being taken down gracefully.
Rt	Number of active routes (prefixes) installed in the routing table.	For inbound RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and outbound RSVP sessions, the routing table is the primary MPLS table (mpls.0).
Active Path	Name of the active path: Primary or Secondary .	This field is used for inbound LSPs only.
P	An asterisk (*) in this column indicates that the LSP is a primary path.	This field is used for inbound LSPs only.
LSPname	Configured name of the LSP.	–
Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	This field is used for outbound and transit LSPs only.
Labelin	Incoming label for this LSP.	–
Labelout	Outgoing label for this LSP.	–
Total	Total number of LSPs displayed for the particular type— ingress (inbound), egress (outbound), or transit .	–

- See Also**
- [Monitoring MPLS Interfaces on page 428](#)
 - [Monitoring MPLS LSP Statistics on page 430](#)
 - [Monitoring RSVP Session Information on page 431](#)

- [Monitoring MPLS RSVP Interfaces Information on page 432](#)

Monitoring MPLS LSP Statistics

Purpose Display statistics for LSP sessions currently active on the device, including the total number of packets and bytes forwarded through an LSP.

Action Select **Monitor>MPLS>LSP Statistics** in the J-Web user interface, or enter the **show mpls lsp statistics** command.



NOTE: Statistics are not available for LSPs on the outbound device, because the penultimate device in the LSP sets the label to 0. Also, as the packet arrives at the outbound device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.

[Table 72 on page 430](#) summarizes key output fields in the MPLS LSP statistics display.

Table 72: Summary of Key MPLS LSP Statistics Output Fields

Field	Values	Additional Information
Ingress LSP	Information about LSPs on the inbound device. Each session has one line of output.	—
Egress LSP	Information about the LSPs on the outbound device. Each session has one line of output.	MPLS learns this information by querying RSVP, which holds all the transit and outbound session information.
Transit LSP	Number of LSPs on the transit routers and the state of these paths.	MPLS learns this information by querying RSVP, which holds all the transit and outbound session information.
To	Destination (outbound device) of the session.	—
From	Source (inbound device) of the session.	—
State	State of the path: Up , Down , or AdminDn .	AdminDn indicates that the LSP is being taken down gracefully.
Packets	Total number of packets received on the LSP from the upstream neighbor.	—
Bytes	Total number of bytes received on the LSP from the upstream neighbor.	—
LSPname	Configured name of the LSP.	—

Table 72: Summary of Key MPLS LSP Statistics Output Fields (*continued*)

Field	Values	Additional Information
Total	Total number of LSPs displayed for the particular type— ingress (inbound), egress (outbound), or transit .	–

- See Also**
- [Monitoring MPLS Interfaces on page 428](#)
 - [Monitoring MPLS LSP Information on page 428](#)
 - [Monitoring RSVP Session Information on page 431](#)
 - [Monitoring MPLS RSVP Interfaces Information on page 432](#)

Monitoring RSVP Session Information

Purpose View information about RSVP-signaled LSP sessions currently active on the device, including inbound (ingress) and outbound (egress) addresses, LSP state, and LSP name.

Action Select **Monitor>MPLS>RSVP Sessions** in the J-Web user interface, or enter the **show rsvp session** command.

[Table 73 on page 431](#) summarizes key output fields in the RSVP session information display.

Table 73: Summary of Key RSVP Session Information Output Fields

Field	Values	Additional Information
Ingress LSP	Information about inbound RSVP sessions. Each session has one line of output.	–
Egress LSP	Information about outbound RSVP sessions. Each session has one line of output.	MPLS learns this information by querying RSVP, which holds all the transit and outbound session information.
Transit LSP	Information about transit RSVP sessions.	MPLS learns this information by querying RSVP, which holds all the transit and outbound session information.
To	Destination (outbound device) of the session.	–
From	Source (inbound device) of the session.	–
State	State of the path: Up , Down , or AdminDn .	AdminDn indicates that the LSP is being taken down gracefully.

Table 73: Summary of Key RSVP Session Information Output Fields (*continued*)

Field	Values	Additional Information
Rt	Number of active routes (prefixes) installed in the routing table.	For inbound RSVP sessions, the routing table is the primary IPv4 table (inet.0). For transit and outbound RSVP sessions, the routing table is the primary MPLS table (mpls.0).
Style	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be FF (fixed filter), SE (shared explicit), or WF (wildcard filter).	This field is used for outbound and transit LSPs only.
Labelin	Incoming label for this RSVP session.	—
Labelout	Outgoing label for this RSVP session.	—
LSPname	Configured name of the LSP.	—
Total	Total number of RSVP sessions displayed for the particular type— ingress (inbound), egress (outbound), or transit .	—

- See Also**
- [Monitoring MPLS Interfaces on page 428](#)
 - [Monitoring MPLS LSP Information on page 428](#)
 - [Monitoring MPLS LSP Statistics on page 430](#)
 - [Monitoring MPLS RSVP Interfaces Information on page 432](#)

Monitoring MPLS RSVP Interfaces Information

Purpose View information about the interfaces on which RSVP is enabled, including the interface name, total bandwidth through the interface, and total current reserved and reservable (available) bandwidth on the interface.

Action Select **Monitor>MPLS>RSVP Interfaces** in the J-Web user interface, or enter the **show rsvp interface** command.

[Table 74 on page 432](#) summarizes key output fields in the RSVP interfaces information display.

Table 74: Summary of Key RSVP Interfaces Information Output Fields

Field	Values	Additional Information
RSVP Interface	Number of interfaces on which RSVP is active. Each interface has one line of output.	—

Table 74: Summary of Key RSVP Interfaces Information Output Fields (*continued*)

Field	Values	Additional Information
Interface	Name of the interface.	—
State	State of the interface: <ul style="list-style-type: none"> • Disabled—No traffic engineering information is displayed. • Down—The interface is not operational. • Enabled—Displays traffic engineering information. • Up—The interface is operational. 	—
Active resv	Number of reservations that are actively reserving bandwidth on the interface.	—
Subscription	User-configured subscription factor.	—
Static BW	Total interface bandwidth, in bits per second (bps).	—
Available BW	Amount of bandwidth that RSVP is allowed to reserve, in bits per second (bps). It is equal to (static bandwidth X subscription factor).	—
Reserved BW	Currently reserved bandwidth, in bits per second (bps).	—
Highwater mark	Highest bandwidth that has ever been reserved on this interface, in bits per second (bps).	—

- See Also**
- [Monitoring MPLS Interfaces on page 428](#)
 - [Monitoring MPLS LSP Information on page 428](#)
 - [Monitoring MPLS LSP Statistics on page 430](#)
 - [Monitoring RSVP Session Information on page 431](#)

- Related Documentation**
- [Understanding Ping MPLS on page 549](#)
 - [MPLS Connection Checking Overview on page 547](#)
 - [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

Monitoring PPP

Supported Platforms SRX1500, SRX300, SRX320, SRX340

Purpose Display PPP monitoring information, including PPP address pool information, session status for PPP interfaces, cumulative statistics for all PPP interfaces, and a summary of PPP sessions.



NOTE: PPP monitoring information is available only in the CLI. The J-Web user interface does not include pages for displaying PPP monitoring information.

Action Enter the following CLI commands:

- `show ppp address-pool pool-name`
- `show ppp interface interface-name`
- `show ppp statistics`
- `show ppp summary`

Related Documentation

- [Monitoring Overview on page 7](#)
- [Monitoring Interfaces on page 426](#)

Monitoring PPPoE

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#)

Purpose Display the session status for PPPoE interfaces, cumulative statistics for all PPPoE interfaces on the device, and the PPPoE version configured on the device.

Action Select **Monitor>PPPoE** in the J-Web user interface. To view interface-specific properties in the J-Web interface, select the interface name on the PPPoE page.

[Table 75 on page 434](#) summarizes key output fields in PPPoE displays.

Table 75: Summary of Key PPPoE Output Fields

Field	Values	Additional Information
PPPoE Interfaces		
Interface	Name of the PPPoE interface.	Click the interface name to display PPPoE information for the interface.
State	State of the PPPoE session on the interface.	—

Table 75: Summary of Key PPPoE Output Fields (*continued*)

Field	Values	Additional Information
Session ID	Unique session identifier for the PPPoE session.	To establish a PPPoE session, first the device acting as a PPPoE client obtains the Ethernet address of the PPPoE server or access concentrator, and then the client and the server negotiate a unique session ID. This process is referred to as PPPoE active discovery and is made up of four steps: initiation, offer, request, and session confirmation. The access concentrator generates the session ID for session confirmation and sends it to the PPPoE client in a PPPoE Active Discovery Session-Confirmation (PADS) packet.
Service Name	Type of service required from the access concentrator.	Service Name identifies the type of service provided by the access concentrator, such as the name of the Internet service provider (ISP), class, or quality of service.
Configured AC Name	Configured access concentrator name.	—
Session AC Names	Name of the access concentrator.	—
AC MAC Address	Media access control (MAC) address of the access concentrator.	—
Session Uptime	Number of seconds the current PPPoE session has been running.	—
Auto-Reconnect Timeout	Number of seconds to wait before reconnecting after a PPPoE session is terminated.	—
Idle Timeout	Number of seconds a PPPoE session can be idle without disconnecting.	—
Underlying Interface	Name of the underlying logical Ethernet or ATM interface on which PPPoE is running—for example, ge-0/0/0.1 .	—
PPPoE Statistics		
Active PPPoE Sessions	Total number of active PPPoE sessions.	—

Table 75: Summary of Key PPPoE Output Fields (*continued*)

Field	Values	Additional Information
Packet Type	Packets sent and received during the PPPoE session, categorized by packet type and packet error: <ul style="list-style-type: none"> • PADI—PPPoE Active Discovery Initiation packets. • PADO—PPPoE Active Discovery Offer packets. • PADR—PPPoE Active Discovery Request packets. • PADS—PPPoE Active Discovery Session-Confirmation packets. • PADT—PPPoE Active Discovery Terminate packets. • Service Name Error—Packets for which the Service-Name request could not be honored. • AC System Error—Packets for which the access concentrator experienced an error in processing the host request. For example, the host had insufficient resources to create a virtual circuit. • Generic Error—Packets that indicate an unrecoverable error occurred. • Malformed Packet—Malformed or short packets that caused the packet handler to disregard the frame as unreadable. • Unknown Packet—Unrecognized packets. 	—
Sent	Number of the specific type of packet sent from the PPPoE client.	—
Received	Number of the specific type of packet received by the PPPoE client.	—
Timeout	Information about the timeouts that occurred during the PPPoE session. <ul style="list-style-type: none"> • PADI—Number of timeouts that occurred for the PADI packet. • PADO—Number of timeouts that occurred for the PADO packet. (This value is always 0 and is not supported.) • PADR—Number of timeouts that occurred for the PADR packet. 	—
Sent	Number of the timeouts that occurred for PADI, PADO, and PADR packets.	—
PPPoE Version		
Maximum Sessions	Maximum number of active PPPoE sessions the device can support. The default is 256 sessions.	—

Table 75: Summary of Key PPPoE Output Fields (*continued*)

Field	Values	Additional Information
PADI Resend Timeout	Initial time, (in seconds) the device waits to receive a PADO packet for the PADI packet sent—for example, 2 seconds. This timeout doubles for each successive PADI packet sent.	The PPPoE Active Discovery Initiation (PADI) packet is sent to the access concentrator to initiate a PPPoE session. Typically, the access concentrator responds to a PADI packet with a PPPoE Active Discovery Offer (PADO) packet. If the access concentrator does not send a PADO packet, the device sends the PADI packet again after timeout period is elapsed. The PADI Resend Timeout doubles for each successive PADI packet sent. For example, if the PADI Resend Timeout is 2 seconds, the second PADI packet is sent after 2 seconds, the third after 4 seconds, the fourth after 8 seconds, and so on.
PADR Resend Timeout	Initial time (in seconds) the device waits to receive a PADS packet for the PADR packet sent. This timeout doubles for each successive PADR packet sent.	The PPPoE Active Discovery Request (PADR) packet is sent to the access concentrator in response to a PADO packet, and to obtain the PPPoE session ID. Typically, the access concentrator responds to a PADR packet with a PPPoE Active Discovery Session-Confirmation (PADS) packet, which contains the session ID. If the access concentrator does not send a PADS packet, the device sends the PADR packet again after the PADR Resend Timeout period is elapsed. The PADR Resend Timeout doubles for each successive PADR packet sent.
Maximum Resend Timeout	Maximum value (in seconds) that the PADI or PADR resend timer can accept—for example, 64 seconds. The maximum value is 64.	—
Maximum Configured AC Timeout	Time (in seconds), within which the configured access concentrator must respond.	—

Alternatively, enter the following CLI commands:

- **show pppoe interfaces**
- **show pppoe statistics**
- **show pppoe version**

You can also view status information about the PPPoE interface by entering the **show interfaces pp0** command in the CLI editor.

Related Documentation

- [Monitoring Overview on page 7](#)
- [Monitoring Interfaces on page 426](#)
- [Monitoring DHCP Client Bindings on page 479](#)

Monitoring Spanning Tree

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

Purpose Use the monitoring functionality to view the Spanning Tree page.

Action To monitor spanning tree, select **Monitor>Switching>Spanning Tree** in the J-Web user interface.

Meaning [Table 76 on page 438](#) summarizes key output fields in the spanning tree page.

Table 76: Spanning Tree Monitoring Page

Field	Value	Additional Information
Bridge parameters		
Context ID	An internally generated identifier.	—
Enabled Protocol	Spanning tree protocol type enabled.	—
Root ID	Bridge ID of the elected spanning tree root bridge.	The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Bridge ID	Locally configured bridge ID.	—
Inter instance ID	An internally generated instance identifier.	—
Extended system ID	Extended system generated instance identifier.	—
Maximum age	Maximum age of received bridge protocol data units (BPDUs).	—
Number of topology changes	Total number of STP topology changes detected since the switch last booted.	—
Forward delay	Spanning tree forward delay.	—
Interface List		
Interface Name	Interface configured to participate in the STP instance.	—
Port ID	Logical interface identifier configured to participate in the STP instance.	—
Designated Port ID	Port ID of the designated port for the LAN segment to which the interface is attached.	—
Port Cost	Configured cost for the interface.	—

Table 76: Spanning Tree Monitoring Page (*continued*)

Field	Value	Additional Information
State	STP port state. Forwarding (FWD), blocking (BLK), listening, learning, or disabled.	–
Role	MSTP or RSTP port role. Designated (DESG), backup (BKUP), alternate (ALT), or root.	–

- Related Documentation**
- [Monitoring Ethernet Switching on page 424](#)
 - [Monitoring GVRP on page 425](#)

Monitoring the WAN Acceleration Interface

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

Purpose View status information and traffic statistics for the WAN acceleration interface.

Action Select **Monitor>WAN Acceleration** in the J-Web user interface, or select **Monitor>Interfaces** and select the interface name (**wx-slot/0/0**). Alternatively, enter the following CLI command:

```
[edit]
user@host# show interfaces wx-slot/0/0 detail
```

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

CHAPTER 25

Monitoring NAT

- [Monitoring NAT on page 441](#)

Monitoring NAT

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring Source NAT Information on page 441](#)
- [Monitoring Destination NAT Information on page 447](#)
- [Monitoring Static NAT Information on page 449](#)
- [Monitoring Incoming Table Information on page 451](#)
- [Monitoring Interface NAT Port Information on page 452](#)

Monitoring Source NAT Information

Purpose Display configured information about source Network Address Translation (NAT) rules, pools, persistent NAT, and paired addresses.

Action Select **Monitor>NAT>Source NAT** in the J-Web user interface, or enter the following CLI commands:

- **show security nat source summary**
- **show security nat source pool *pool-name***
- **show security nat source persistent-nat-table**
- **show security nat source paired-address**

[Table 77 on page 441](#) describes the available options for monitoring source NAT.

Table 77: Source NAT Monitoring Page

Field	Description	Action
Rules		
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.

Table 77: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Total rules	Number of rules configured.	—
ID	Rule ID number.	—
Name	Name of the rule .	—
From	Name of the routing instance/zone/interface from which the packet flows.	—
To	Name of the routing instance/zone/interface to which the packet flows.	—
Source address range	Source IP address range in the source pool.	—
Destination address range	Destination IP address range in the source pool.	—
Source ports	Source port numbers.	—
Ip protocol	IP protocol.	—
Action	Action taken for a packet that matches a rule.	—
Persistent NAT type	Persistent NAT type.	—
Inactivity timeout	Inactivity timeout interval for the persistent NAT binding.	—
Alarm threshold	Utilization alarm threshold.	
Max session number	The maximum number of sessions.	—
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. <ul style="list-style-type: none"> • Succ—Number of successful session installations after the NAT rule is matched. • Failed—Number of unsuccessful session installations after the NAT rule is matched. • Current—Number of sessions that reference the specified rule. 	—

Table 77: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Translation Hits	Number of times a translation in the translation table is used for a source NAT rule.	—
Pools		
Pool Name	The names of the pools.	Select all pools or a specific pool to display from the list.
Total Pools	Total pools added.	—
ID	ID of the pool.	—
Name	Name of the source pool.	—
Address range	IP address range in the source pool.	—
Single/Twin ports	Number of allocated single and twin ports.	—
Port	Source port number in the pool.	—
Address assignment	Displays the type of address assignment.	—
Alarm threshold	Utilization alarm threshold.	—
Port overloading factor	Port overloading capacity.	—
Routing instance	Name of the routing instance.	—
Total addresses	Total IP address, IP address set, or address book entry.	—
Host address base	Host base address of the original source IP address range.	—
Translation hits	Number of times a translation in the translation table is used for source NAT.	—
Top 10 Translation Hits		
Graph	Displays the graph of top 10 translation hits.	—
Persistent NAT		

Table 77: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Persistent NAT table statistics		
binding total	Displays the total number of persistent NAT bindings for the FPC.	—
binding in use	Number of persistent NAT bindings that are in use for the FPC.	—
enode total	Total number of persistent NAT enodes for the FPC.	—
enode in use	Number of persistent NAT enodes that are in use for the FPC.	—
Persistent NAT table		
Source NAT pool	Name of the pool.	Select all pools or a specific pool to display from the list.
Internal IP	Internal IP address.	Select all IP addresses or a specific IP address to display from the list.
Internal port	Displays the internal ports configured in the system.	Select the port to display from the list.
Internal protocol	Internal protocols .	Select all protocols or a specific protocol to display from the list.
Internal IP	Internal transport IP address of the outgoing session from internal to external.	—
Internal port	Internal transport port number of the outgoing session from internal to external.	—
Internal protocol	Internal protocol of the outgoing session from internal to external.	—
Reflective IP	Translated IP address of the source IP address.	—
Reflective port	Displays the translated number of the port.	—
Reflective protocol	Translated protocol.	—
Source NAT pool	Name of the source NAT pool where persistent NAT is used.	—
Type	Persistent NAT type.	—

Table 77: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Left time/Conf time	Inactivity timeout period that remains and the configured timeout value.	—
Current session num/Max session num	Number of current sessions associated with the persistent NAT binding and the maximum number of sessions.	—
Source NAT rule	Name of the source NAT rule to which this persistent NAT binding applies.	—
External node table		
Internal IP	Internal transport IP address of the outgoing session from internal to external.	—
Internal port	Internal port number of the outgoing session from internal to external.	—
External IP	External IP address of the outgoing session from internal to external.	—
External port	External port of the outgoing session from internal to external.	—
Zone	External zone of the outgoing session from internal to external.	—
Paired Address		
Pool name	Name of the pool.	Select all pools or a specific pool to display from the list.
Specified Address	IP address.	Select all addresses, or select the internal or external IP address to display, and enter the IP address.
Pool name	Displays the selected pool or pools.	—
Internal address	Displays the internal IP address.	—
External address	Displays the external IP address.	—
Resource Usage		
Utilization for all source pools		

Table 77: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Pool name	Name of the pool.	To view additional usage information for Port Address Translation (PAT) pools, select a pool name. The information displays under Detail Port Utilization for Specified Pool.
Pool type	Pool type: PAT or Non-PAT.	—
Port overloading factor	Port overloading capacity for PAT pools.	—
Address	Addresses in the pool.	—
Used	<p>Number of used resources in the pool.</p> <p>For Non-PAT pools, the number of used IP addresses is displayed.</p> <p>For PAT pools, the number of used ports is displayed.</p>	—
Available	<p>Number of available resources in the pool.</p> <p>For Non-PAT pools, the number of available IP addresses is displayed.</p> <p>For PAT pools, the number of available ports is displayed.</p>	—
Total	<p>Number of used and available resources in the pool.</p> <p>For Non-PAT pools, the total number of used and available IP addresses is displayed.</p> <p>For PAT pools, the total number of used and available ports is displayed.</p>	—
Usage	<p>Percent of resources used.</p> <p>For Non-PAT pools, the percent of IP addresses used is displayed.</p> <p>For PAT pools, the percent of ports, including single and twin ports, is displayed.</p>	—
Peak usage	Percent of resources used during the peak date and time.	—
Detail Port Utilization for Specified Pool		
Address Name	IP addresses in the PAT pool.	Select the IP address for which you want to display detailed usage information.
Factor-Index	Index number.	—

Table 77: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Port-range	Displays the number of ports allocated at a time.	—
Used	Displays the number of used ports.	—
Available	Displays the number of available ports.	—
Total	Displays the number of used and available ports.	—
Usage	Displays the percentage of ports used during the peak date and time.	—

- See Also**
- [Monitoring Destination NAT Information on page 447](#)
 - [Monitoring Static NAT Information on page 449](#)
 - [Monitoring Incoming Table Information on page 451](#)
 - [Monitoring Interface NAT Port Information on page 452](#)

Monitoring Destination NAT Information

Purpose View the destination Network Address Translation (NAT) summary table and the details of the specified NAT destination address pool information.

Action Select **Monitor>NAT> Destination NAT** in the J-Web user interface, or enter the following CLI commands:

- **show security nat destination summary**
- **show security nat destination pool *pool-name***

[Table 78 on page 447](#) summarizes key output fields in the destination NAT display.

Table 78: Summary of Key Destination NAT Output Fields

Field	Values	Action
Rules		
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	—
ID	Rule ID number.	—
Name	Name of the rule .	—

Table 78: Summary of Key Destination NAT Output Fields (*continued*)

Field	Values	Action
Ruleset Name	Name of the rule set.	—
From	Name of the routing instance/zone/interface from which the packet flows.	—
Source address range	Source IP address range in the source pool.	—
Destination address range	Destination IP address range in the source pool.	—
Destination port	Destination port in the destination pool.	—
IP protocol	IP protocol.	—
Action	Action taken for a packet that matches a rule.	—
Alarm threshold	Utilization alarm threshold.	—
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. <ul style="list-style-type: none"> • Succ—Number of successful session installations after the NAT rule is matched. • Failed—Number of unsuccessful session installations after the NAT rule is matched. • Current—Number of sessions that reference the specified rule. 	—
Translation hits	Number of times a translation in the translation table is used for a destination NAT rule.	—
Pools		
Pool Name	The names of the pools.	Select all pools or a specific pool to display from the list.
Total Pools	Total pools added.	—
ID	ID of the pool.	—
Name	Name of the destination pool.	—
Address range	IP address range in the destination pool.	—

Table 78: Summary of Key Destination NAT Output Fields (*continued*)

Field	Values	Action
Port	Destination port number in the pool.	—
Routing instance	Name of the routing instance.	—
Total addresses	Total IP address, IP address set, or address book entry.	—
Translation hits	Number of times a translation in the translation table is used for destination NAT.	—
Top 10 Translation Hits		
Graph	Displays the graph of top 10 translation hits.	—

- See Also**
- [Monitoring Source NAT Information on page 441](#)
 - [Monitoring Static NAT Information on page 449](#)
 - [Monitoring Incoming Table Information on page 451](#)
 - [Monitoring Interface NAT Port Information on page 452](#)

Monitoring Static NAT Information

Purpose View static NAT rule information.

Action Select **Monitor>NAT>Static NAT** in the J-Web user interface, or enter the following CLI command:

```
show security nat static rule
```

[Table 79 on page 449](#) summarizes key output fields in the static NAT display.

Table 79: Summary of Key Static NAT Output Fields

Field	Values	Action
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	—
ID	Rule ID number.	—
Position	Position of the rule that indicates the order in which it applies to traffic.	—

Table 79: Summary of Key Static NAT Output Fields (*continued*)

Field	Values	Action
Name	Name of the rule.	–
Ruleset Name	Name of the rule set.	–
From	Name of the routing instance/interface/zone from which the packet comes	–
Source addresses	Source IP addresses.	–
Source ports	Source port numbers.	–
Destination addresses	Destination IP address and subnet mask.	–
Destination ports	Destination port numbers .	–
Host addresses	Name of the host addresses.	–
Host ports	Host port numbers.	–
Netmask	Subnet IP address.	–
Host routing instance	Name of the routing instance from which the packet comes.	–
Alarm threshold	Utilization alarm threshold.	–
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. <ul style="list-style-type: none"> • Succ—Number of successful session installations after the NAT rule is matched. • Failed—Number of unsuccessful session installations after the NAT rule is matched. • Current—Number of sessions that reference the specified rule. 	–
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.	–
Top 10 Translation Hits Graph	Displays the graph of top 10 translation hits.	–

- See Also**
- [Monitoring Source NAT Information on page 441](#)
 - [Monitoring Destination NAT Information on page 447](#)

- [Monitoring Incoming Table Information on page 451](#)
- [Monitoring Interface NAT Port Information on page 452](#)

Monitoring Incoming Table Information

Purpose View NAT table information.

Action Select **Monitor>NAT>Incoming Table** in the J-Web user interface, or enter the following CLI command:

show security nat incoming-table

[Table 80 on page 451](#) summarizes key output fields in the incoming table display.

Table 80: Summary of Key Incoming Table Output Fields

Field	Values
Statistics	
In use	Number of entries in the NAT table.
Maximum	Maximum number of entries possible in the NAT table.
Entry allocation failed	Number of entries failed for allocation.
Incoming Table	
Clear	
Destination	Destination IP address and port number.
Host	Host IP address and port number that the destination IP address is mapped to.
References	Number of sessions referencing the entry.
Timeout	Timeout, in seconds, of the entry in the NAT table.
Source-pool	Name of source pool where translation is allocated.

- See Also**
- [Monitoring Source NAT Information on page 441](#)
 - [Monitoring Destination NAT Information on page 447](#)
 - [Monitoring Static NAT Information on page 449](#)
 - [Monitoring Interface NAT Port Information on page 452](#)

Monitoring Interface NAT Port Information

Purpose View port usage for an interface source pool information.

Action To monitoring interface NAT port information, do one of the following:

- If you are using SRX5400, SRX5600, or SRX5800 platforms, select **Monitor>Firewall/NAT>Interface NAT** in the J-Web user interface or enter the CLI command **show security nat interface-nat-ports**.
- Select **Monitor>NAT>Interface NAT Ports** in the J-Web user interface.

Table 81 on page 452 summarizes key output fields in the interface NAT display.

Table 81: Summary of Key Interface NAT Output Fields

Field	Values	Additional Information
Interface NAT Summary Table		
Pool Index	Port pool index.	—
Total Ports	Total number of ports in a port pool.	—
Single Ports Allocated	Number of ports allocated one at a time that are in use.	—
Single Ports Available	Number of ports allocated one at a time that are free for use.	—
Twin Ports Allocated	Number of ports allocated two at a time that are in use.	—
Twin Ports Available	Number of ports allocated two at a time that are free for use.	—

- See Also**
- [Monitoring Source NAT Information on page 441](#)
 - [Monitoring Destination NAT Information on page 447](#)
 - [Monitoring Static NAT Information on page 449](#)
 - [Monitoring Incoming Table Information on page 451](#)

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

Monitoring Security Policies

- [Monitoring Policy Statistics on page 453](#)
- [Monitoring Routing Information on page 454](#)
- [Monitoring Security Events by Policy on page 462](#)
- [Monitoring Security Features on page 464](#)

Monitoring Policy Statistics

Supported Platforms [SRX Series, vSRX](#)

Purpose Monitor and record traffic that Junos OS permits or denies based on previously configured policies.

Action To monitor traffic, enable the count and log options.

Count—Configurable in an individual policy. If count is enabled, statistics are collected for sessions that enter the device for a given policy, and for the number of packets and bytes that pass through the device in both directions for a given policy. For counts (only for packets and bytes), you can specify that alarms be generated whenever the traffic exceeds specified thresholds. See [count \(Security Policies\)](#).

Log—Logging capability can be enabled with security policies during session initialization (**session-init**) or session close (**session-close**) stage. See [log \(Security Policies\)](#).

- To view logs from denied connections, enable log on **session-init**.
- To log sessions after their conclusion/tear-down, enable log on **session-close**.



NOTE: Session log is enabled at real time in the flow code which impacts the user performance. If both **session-close** and **session-init** are enabled, performance is further degraded as compared to enabling **session-init** only.

For details about information collected for session logs, see *Information Provided in Session Log Entries for SRX Series Services Gateways*.

- Related Documentation**
- [Security Policies Overview](#)
 - [Troubleshooting Security Policies on page 604](#)
 - [Checking a Security Policy Commit Failure on page 604](#)
 - [Verifying a Security Policy Commit on page 605](#)
 - [Debugging Policy Lookup on page 606](#)

Monitoring Routing Information

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring Route Information on page 454](#)
- [Monitoring RIP Routing Information on page 456](#)
- [Monitoring OSPF Routing Information on page 458](#)
- [Monitoring BGP Routing Information on page 460](#)

Monitoring Route Information

Purpose View information about the routes in a routing table, including destination, protocol, state, and parameter information.

Action Select **Monitor>Routing>Route Information** in the J-Web user interface, or enter the following CLI commands:

- `show route terse`
- `show route detail`



NOTE: When you use an HTTPS connection in the Microsoft Internet Explorer browser to save a report from this page in the J-Web interface, the error message "Internet Explorer was not able to open the Internet site" is displayed. This problem occurs because the Cache-Control: no cache HTTP header is added on the server side and Internet Explorer does not allow you to download the encrypted file with the Cache-Control: no cache HTTP header set in the response from the server.

As a workaround, refer to Microsoft Knowledge Base article 323308, which is available at this URL: <http://support.microsoft.com/kb/323308>. Also, you can alternatively use HTTP in the Internet Explorer browser or use HTTPS in the Mozilla Firefox browser to save a file from this page.

[Table 82 on page 455](#) describes the different filters, their functions, and the associated actions.

Table 83 on page 455 summarizes key output fields in the routing information display.

Table 82: Filtering Route Messages

Field	Function	Your Action
Destination Address	Specifies the destination address of the route.	Enter the destination address.
Protocol	Specifies the protocol from which the route was learned.	Enter the protocol name.
Next hop address	Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	Enter the next hop address.
Receive protocol	Specifies the dynamic routing protocol using which the routing information was received through a particular neighbor.	Enter the routing protocol.
Best route	Specifies only the best route available.	Select the view details of the best route.
Inactive routes	Specifies the inactive routes.	Select the view details of inactive routes.
Exact route	Specifies the exact route.	Select the view details of the exact route.
Hidden routes	Specifies the hidden routes.	Select the view details of hidden routes.
Search	Applies the specified filter and displays the matching messages.	To apply the filter and display messages, click Search .
Reset	Resets selected options to default	To reset the filter, click Reset .

Table 83: Summary of Key Routing Information Output Fields

Field	Values	Additional Information
Static Route Addresses	The list of static route addresses.	—
Protocol	Protocol from which the route was learned: Static , Direct , Local , or the name of a particular protocol.	—
Preference	The preference is the individual preference value for the route.	The route preference is used as one of the route selection criteria.

Table 83: Summary of Key Routing Information Output Fields (*continued*)

Field	Values	Additional Information
Next-Hop	Network Layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	<p>If a next hop is listed as Discard, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the discard attribute has been set.</p> <p>If a next hop is listed as Reject, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected.</p> <p>If a next hop is listed as Local, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).</p>
Age	How long the route has been active.	—
State	Flags for this route.	There are many possible flags.
AS Path	<p>AS path through which the route was learned. The letters of the AS path indicate the path origin:</p> <ul style="list-style-type: none"> • I—IGP. • E—EGP. • ?—Incomplete. Typically, the AS path was aggregated. 	—

- See Also**
- [Monitoring RIP Routing Information on page 456](#)
 - [Monitoring OSPF Routing Information on page 458](#)
 - [Monitoring BGP Routing Information on page 460](#)

Monitoring RIP Routing Information

Purpose View RIP routing information, including a summary of RIP neighbors and statistics.

Action Select **Monitor>Routing>RIP Information** in the J-Web user interface, or enter the following CLI commands:

- **show rip statistics**
- **show rip neighbors**

[Table 84 on page 457](#) summarizes key output fields in the RIP routing display in the J-Web user interface.

Table 84: Summary of Key RIP Routing Output Fields

Field	Values	Additional Information
RIP Statistics		
Protocol Name	The RIP protocol name.	—
Port number	The port on which RIP is enabled.	—
Hold down time	The interval during which routes are neither advertised nor updated.	—
Global routes learned	Number of RIP routes learned on the logical interface.	—
Global routes held down	Number of RIP routes that are not advertised or updated during the hold-down interval.	—
Global request dropped	Number of requests dropped.	—
Global responses dropped	Number of responses dropped.	—
RIP Neighbors		
Details	Tab used to view the details of the interface on which RIP is enabled.	—
Neighbor	Name of the RIP neighbor.	This value is the name of the interface on which RIP is enabled. Click the name to see the details for this neighbor.
State	State of the RIP connection: Up or Dn (Down).	—
Source Address	Local source address.	This value is the configured address of the interface on which RIP is enabled.
Destination Address	Destination address.	This value is the configured address of the immediate RIP adjacency.
Send Mode	The mode of sending RIP messages.	—
Receive Mode	The mode in which messages are received.	—
In Metric	Value of the incoming metric configured for the RIP neighbor.	—

- See Also**
- [Monitoring Route Information on page 454](#)
 - [Monitoring OSPF Routing Information on page 458](#)

- [Monitoring BGP Routing Information on page 460](#)

Monitoring OSPF Routing Information

Purpose View OSPF routing information, including a summary of OSPF neighbors, interfaces, and statistics.

Action Select **Monitor>Routing>OSPF Information** in the J-Web user interface, or enter the following CLI commands:

- **show ospf neighbors**
- **show ospf interfaces**
- **show ospf statistics**

[Table 85 on page 458](#) summarizes key output fields in the OSPF routing display in the J-Web user interface.

Table 85: Summary of Key OSPF Routing Output Fields

Field	Values	Additional Information
OSPF Interfaces		
Details	Tab used to view the details of the selected OSPF.	—
Interface	Name of the interface running OSPF.	—
State	State of the interface: BDR , Down , DR , DRother , Loop , PtToPt , or Waiting .	The Down state, indicating that the interface is not functioning, and PtToPt state, indicating that a point-to-point connection has been established, are the most common states.
Area	Number of the area that the interface is in.	—
DR ID	ID of the area's designated device.	—
BDR ID	ID of the area's backup designated device.	—
Neighbors	Number of neighbors on this interface.	—
OSPF Statistics		
Packets tab		
Sent	Displays the total number of packets sent.	—
Received	Displays the total number of packets received.	—
Details tab		
Flood Queue Depth	Number of entries in the extended queue.	—

Table 85: Summary of Key OSPF Routing Output Fields (*continued*)

Field	Values	Additional Information
Total Retransmits	Number of retransmission entries enqueued.	–
Total Database Summaries	Total number of database description packets.	–
OSPF Neighbors		
Address	Address of the neighbor.	–
Interface	Interface through which the neighbor is reachable.	–
State	State of the neighbor: Attempt, Down, Exchange, ExStart, Full, Init, Loading, or 2way.	Generally, only the Down state, indicating a failed OSPF adjacency, and the Full state, indicating a functional adjacency, are maintained for more than a few seconds. The other states are transitional states that a neighbor is in only briefly while an OSPF adjacency is being established.
ID	ID of the neighbor.	–
Priority	Priority of the neighbor to become the designated router.	–
Activity Time	The activity time.	–
Area	Area that the neighbor is in.	–
Options	Option bits received in the hello packets from the neighbor.	–
DR Address	Address of the designated router.	–
BDR Address	Address of the backup designated router.	–
Uptime	Length of time since the neighbor came up.	–
Adjacency	Length of time since the adjacency with the neighbor was established.	–

- See Also**
- [Monitoring Route Information on page 454](#)
 - [Monitoring RIP Routing Information on page 456](#)
 - [Monitoring BGP Routing Information on page 460](#)

Monitoring BGP Routing Information

Purpose Monitor BGP routing information on the routing device, including a summary of BGP routing and neighbor information.

Action Select **Monitor>Routing>BGP Information** in the J-Web user interface, or enter the following CLI commands:

- **show bgp summary**
- **show bgp neighbor**

Table 86 on page 460 summarizes key output fields in the BGP routing display in the J-Web user interface.

Table 86: Summary of Key BGP Routing Output Fields

Field	Values	Additional Information
BGP Peer Summary		
Total Groups	Number of BGP groups.	—
Total Peers	Number of BGP peers.	—
Down Peers	Number of unavailable BGP peers.	—
Unconfigured Peers	Address of each BGP peer.	—
RIB Summary tab		
RIB Name	Name of the RIB group.	—
Total Prefixes	Total number of prefixes from the peer, both active and inactive, that are in the routing table.	—
Active Prefixes	Number of prefixes received from the EBGp peers that are active in the routing table.	—
Suppressed Prefixes	Number of routes received from EBGp peers currently inactive because of damping or other reasons.	—
History Prefixes	History of the routes received or suppressed.	—
Dumped Prefixes	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	—
Pending Prefixes	Number of pending routes.	—

Table 86: Summary of Key BGP Routing Output Fields (*continued*)

Field	Values	Additional Information
State	Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete.	—
BGP Neighbors		
Details	Click this button to view the selected BGP neighbor details.	—
Peer Address	Address of the BGP neighbor.	—
Autonomous System	AS number of the peer.	—
Peer State	<p>Current state of the BGP session:</p> <ul style="list-style-type: none"> • Active—BGP is initiating a TCP connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message. • Connect—BGP is waiting for the TCP connection to become complete. • Established—The BGP session has been established, and the peers are exchanging BGP update messages. • Idle—This is the first stage of a connection. BGP is waiting for a Start event. • OpenConfirm—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • OpenSent—BGP has sent an open message and is waiting to receive an open message from the peer. 	Generally, the most common states are Active , which indicates a problem establishing the BGP connection, and Established , which indicates a successful session setup. The other states are transition states, and BGP sessions normally do not stay in those states for extended periods of time.
Elapsed Time	Elapsed time since the peering session was last reset.	—
Description	Description of the BGP session.	—

- See Also**
- [Monitoring Route Information on page 454](#)
 - [Monitoring RIP Routing Information on page 456](#)
 - [Monitoring OSPF Routing Information on page 458](#)

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

Monitoring Security Events by Policy

Supported Platforms SRX Series, vSRX

Purpose Monitor security events by policy and display logged event details with the J-Web user interface.

Action To monitor security events by policy:

1. Select one of the following in the J-Web user interface:
 - If you are using SRX5400, SRX5600, or SRX5800 platforms, select **Monitor>Events and Alarms>Security Events**.
 - Select **Monitor>Alarms>Policy Log**.

The View Policy Log pane appears. [Table 87 on page 462](#) describes the content of this pane.

Table 87: View Policy Log Fields

Field	Value
Log file name	Name of the event log files to search.
Policy name	Name of the policy of the events to be retrieved.
Source address	Source address of the traffic that triggered the event.
Destination address	Destination address of the traffic that triggered the event.
Event type	Type of event that was triggered by the traffic.
Application	Application of the traffic that triggered the event.
Source port	Source port of the traffic that triggered the event.
Destination port	Destination port of the traffic that triggered the event.
Source zone	Source zone of the traffic that triggered the event.
Destination zone	Destination zone of the traffic that triggered the event.
Source NAT rule	Source NAT rule of the traffic that triggered the event.
Destination NAT rule	Destination NAT rule of the traffic that triggered the event.
Is global policy	Specifies that the policy is a global policy.

If your device is not configured to store session log files locally, the Create log configuration button is displayed in the lower-right portion of the View Policy Log pane.

- To store session log files locally, click **Create log configuration**.

If session logs are being sent to an external log collector (stream mode has been configured for log files), a message appears indicating that event mode must be configured to view policy logs.



NOTE: Reverting to event mode will discontinue event logging to the external log collector.

- To reset the **mode** option to **event**, enter the **set security log** command.
2. Enter one or more search fields in the View Policy Log pane and click **Search** to display events matching your criteria.

For example, enter the event type **Session Close** and the policy **pol1** to display event details from all Session Close logs that contain the specified policy. To reduce search results further, add more criteria about the particular event or group of events that you want displayed.

The Policy Events Detail pane displays information from each matching session log. [Table 88 on page 463](#) describes the contents of this pane.

Table 88: Policy Events Detail Fields

Field	Value
Timestamp	Time when the event occurred.
Policy name	Policy that triggered the event.
Record type	Type of event log providing the data.
Source IP/Port	Source address (and port, if applicable) of the event traffic.
Destination IP/Port	Destination address (and port, if applicable) of the event traffic.
Service name	Service name of the event traffic.
NAT source IP/Port	NAT source address (and port, if applicable) of the event traffic.
NAT destination IP/Port	NAT destination address (and port, if applicable) of the event traffic.

Related Documentation • [Monitoring Overview on page 7](#)

- [Monitoring Interfaces on page 426](#)
- [Monitoring Alarms on page 339](#)
- [Monitoring Events on page 480](#)

Monitoring Security Features

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring Policies on page 464](#)
- [Checking Policies on page 466](#)
- [Monitoring Screen Counters on page 469](#)
- [Monitoring IDP Status on page 472](#)
- [Monitoring Flow Gate Information on page 473](#)
- [Monitoring Firewall Authentication Table on page 474](#)
- [Monitoring Firewall Authentication History on page 476](#)
- [Monitoring 802.1x on page 477](#)

Monitoring Policies

Purpose Display, sort, and review policy activity for every activated policy configured on the device. Policies are grouped by Zone Context (the from and to zones of the traffic) to control the volume of data displayed at one time. From the policy list, select a policy to display statistics and current network activity.

Action To review policy activity:

1. Select **Monitor>Security>Policy>Activities** in the J-Web user interface. The Security Policies Monitoring page appears and lists the policies from the first Zone Context. See [Table 89 on page 464](#) for field descriptions.
2. Select the **Zone Context** of the policy you want to monitor, and click **Filter**. All policies within the zone context appear in match sequence.
3. Select a policy, and click **Clear Statistics** to set all counters to zero for the selected policy.

Table 89: Security Policies Monitoring Output Fields

Field	Value	Additional Information
Zone Context (Total #)	Displays a list of all from and to zone combinations for the configured policies. The total number of active policies for each context is specified in the Total # field. By default, the policies from the first Zone Context are displayed.	To display policies for a different context, select a zone context and click Filter . Both inactive and active policies appear for each context. However, the Total # field for a context specifies the number of active policies only.

Table 89: Security Policies Monitoring Output Fields (*continued*)

Field	Value	Additional Information
Default Policy action	Specifies the action to take for traffic that does not match any of the policies in the context: <ul style="list-style-type: none"> • permit-all—Permit all traffic that does not match a policy. • deny-all—Deny all traffic that does not match a policy. 	—
From Zone	Displays the source zone to be used as match criteria for the policy.	—
To Zone	Displays the destination zone to be used as match criteria for the policy.	—
Name	Displays the name of the policy.	—
Source Address	Displays the source addresses to be used as match criteria for the policy. Address sets are resolved to their individual names. (In this case, only the names are given, not the IP addresses).	—
Destination Address	Displays the destination addresses (or address sets) to be used as match criteria for the policy. Addresses are entered as specified in the destination zone's address book.	—
Source Identity	Displays the name of the source identities set for the policy.	To display the value of the source identities, hover the mouse on this field. Unknown source identities are also displayed.
Application	Displays the name of a predefined or custom application signature to be used as match criteria for the policy.	—
Dynamic App	Displays the dynamic application signatures to be used as match criteria if an application firewall rule set is configured for the policy. For a network firewall, a dynamic application is not defined.	The rule set appears in two lines. The first line displays the configured dynamic application signatures in the rule set. The second line displays the default dynamic application signature. If more than two dynamic application signatures are specified for the rule set, hover over the output field to display the full list in a tooltip.
Action	Displays the action portion of the rule set if an application firewall rule set is configured for the policy. <ul style="list-style-type: none"> • permit—Permits access to the network services controlled by the policy. A green background signifies permission. • deny—Denies access to the network services controlled by the policy. A red background signifies denial. 	The action portion of the rule set appears in two lines. The first line identifies the action to be taken when the traffic matches a dynamic application signature. The second line displays the default action when traffic does not match a dynamic application signature.

Table 89: Security Policies Monitoring Output Fields (*continued*)

Field	Value	Additional Information
NW Services	<p>Displays the network services permitted or denied by the policy if an application firewall rule set is configured. Network services include:</p> <ul style="list-style-type: none"> • gprs-gtp-profile—Specify a GPRS Tunneling Protocol profile name. • idp—Perform intrusion detection and prevention. • redirect-wx—Set WX redirection. • reverse-redirect-wx—Set WX reverse redirection. • uac-policy—Enable unified access control enforcement of the policy. 	—
Policy Hit Counters Graph	Provides a representation of the value over time for a specified counter. The graph is blank if Policy Counters indicates no data. As a selected counter accumulates data, the graph is updated at each refresh interval.	To toggle a graph on and off, click the counter name below the graph.
Policy Counters	<p>Lists statistical counters for the selected policy if Count is enabled. The following counters are available for each policy:</p> <ul style="list-style-type: none"> • input-bytes • input-byte-rate • output-bytes • output-byte-rate • input-packets • input-packet-rate • output-packets • output-packet-rate • session-creations • session-creation-rate • active-sessions 	To graph or to remove a counter from the Policy Hit Counters Graph, toggle the counter name. The names of enabled counters appear below the graph.

- See Also**
- [Checking Policies on page 466](#)
 - [Monitoring Screen Counters on page 469](#)

Checking Policies

Purpose Enter match criteria and conduct a policy search. The search results include all policies that match the traffic criteria in the sequence in which they will be encountered.

Because policy matches are listed in the sequence in which they would be encountered, you can determine whether a specific policy is being applied correctly or not. The first policy in the list is applied to all matching traffic. Policies listed after this one remain in the “shadow” of the first policy and are never encountered by this traffic.

By manipulating the traffic criteria and policy sequence, you can tune policy application to suit your needs. During policy development, you can use this feature to establish the appropriate sequence of policies for optimum traffic matches. When troubleshooting, use this feature to determine if specific traffic is encountering the appropriate policy.

- Action**
1. Select **Monitor>Security>Policy>Shadow Policies** in the J-Web user interface. The Check Policies page appears. [Table 90 on page 467](#) explains the content of this page.
 2. In the top pane, enter the From Zone and To Zone to supply the context for the search.
 3. Enter match criteria for the traffic, including the source address and port, the destination address and port, and the protocol of the traffic.
 4. Enter the number of matching policies to display.
 5. Click **Search** to find policies matching your criteria. The lower pane displays all policies matching the criteria up to the number of policies you specified.
 - The first policy will be applied to all traffic with this match criteria.
 - Remaining policies will not be encountered by any traffic with this match criteria.
 6. To manipulate the position and activation of a policy, select the policy and click the appropriate button:
 - **Move**—Moves the selected policy up or down to position it at a more appropriate point in the search sequence.
 - **Move to**—Moves the selected policy by allowing you to drag and drop it to a different location on the same page.

Table 90: Check Policies Output

Field	Function
Check Policies Search Input Pane	
From Zone	Name or ID of the source zone. If a From Zone is specified by name, the name is translated to its ID internally.
To Zone	Name or ID of the destination zone. If a To Zone is specified by name, the name is translated to its ID internally.
Source Address	Address of the source in IP notation.
Source Port	Port number of the source.
Destination Address	Address of the destination in IP notation.
Destination Port	Port number of the destination.

Table 90: Check Policies Output (*continued*)

Field	Function
Source Identity	Name of the source identity.
Protocol	Name or equivalent value of the protocol to be matched. ah —51 egp —8 esp —50 gre —47 icmp —1 igmp —2 igp —9 ipip —94 ipv6 —41 ospf —89 pgm —113 pim —103 rdp —27 rsvp —46 sctp —132 tcp —6 udp —17 vrrp —112
Result Count	(Optional) Number of policies to display. Default value is 1. Maximum value is 16.
Check Policies List	
From Zone	Name of the source zone.
To Zone	Name of the destination zone.
Total Policies	Number of policies retrieved.
Default Policy action	The action to be taken if no match occurs.

Table 90: Check Policies Output (*continued*)

Field	Function
Name	Policy name
Source Address	Name of the source address (not the IP address) of a policy. Address sets are resolved to their individual names.
Destination Address	Name of the destination address or address set. A packet's destination address must match this value for the policy to apply to it.
Source Identity	Name of the source identity for the policy.
Application	Name of a preconfigured or custom application of the policy match.
Action	Action taken when a match occurs as specified in the policy.
Hit Counts	Number of matches for this policy. This value is the same as the Policy Lookups in a policy statistics report.
Active Sessions	Number of active sessions matching this policy.

Alternatively, to list matching policies using the CLI, enter the **show security match-policies** command and include your match criteria and the number of matching policies to display.

- See Also**
- [Monitoring Policies on page 464](#)
 - [Monitoring Screen Counters on page 469](#)

Monitoring Screen Counters

Purpose View screen statistics for a specified security zone.

Action Select **Monitor>Security>Screen Counters** in the J-Web user interface, or enter the following CLI command:

show security screen statistics zone zone-name

[Table 91 on page 469](#) summarizes key output fields in the screen counters display.

Table 91: Summary of Key Screen Counters Output Fields

Field	Values	Additional Information
Zones		
ICMP Flood	Internet Control Message Protocol (ICMP) flood counter.	An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.

Table 91: Summary of Key Screen Counters Output Fields (*continued*)

Field	Values	Additional Information
UDP Flood	User Datagram Protocol (UDP) flood counter.	UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP Winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks.	WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP Port Scan	Number of TCP port scans.	The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP Address Sweep	Number of ICMP address sweeps.	An IP address sweep can occur with the intent of triggering responses from active hosts.
IP Tear Drop	Number of teardrop attacks.	Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN Attack	Number of TCP SYN attacks.	—
IP Spoofing	Number of IP spoofs.	IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP Ping of Death	ICMP ping of death counter.	Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP Source Route	Number of IP source route attacks.	—
TCP Land Attack	Number of land attacks.	Land attacks occur when attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN Fragment	Number of TCP SYN fragments.	—
TCP No Flag	Number of TCP headers without flags set.	A normal TCP segment header has at least one control flag set.
IP Unknown Protocol	Number of unknown Internet protocols.	—
IP Bad Options	Number of invalid options.	—
IP Record Route Option	Number of packets with the IP record route option enabled.	This option records the IP addresses of the network devices along the path that the IP packet travels.

Table 91: Summary of Key Screen Counters Output Fields (*continued*)

Field	Values	Additional Information
IP Timestamp Option	Number of IP timestamp option attacks.	This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP Security Option	Number of IP security option attacks.	–
IP Loose route Option	Number of IP loose route option attacks.	This option specifies a partial route list for a packet to take on its journey from source to destination.
IP Strict Source Route Option	Number of IP strict source route option attacks.	This option specifies the complete route list for a packet to take on its journey from source to destination.
IP Stream Option	Number of stream option attacks.	This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP Fragment	Number of ICMP fragments.	Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP Large Packet	Number of large ICMP packets.	–
TCP SYN FIN Packet	Number of TCP SYN FIN packets.	–
TCP FIN without ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.	–
TCP SYN-ACK-ACK Proxy	Number of TCP flags enabled with SYN-ACK-ACK.	To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, Junos OS rejects further connection requests from that IP address.
IP Block Fragment	Number of IP block fragments.	–

See Also

- [Monitoring Policies on page 464](#)
- [Checking Policies on page 466](#)

Monitoring IDP Status

Purpose View detailed information about the IDP Status, Memory, Counters, Policy Rulebase Statistics, and Attack table statistics.

Action To view Intrusion Detection and Prevention (IDP) table information, do one of the following:

- If you are using SRX5400, SRX5600, or SRX5800 platforms, select **Monitor>Security>IDP>Status** in the J-Web user interface, or enter the following CLI commands:

show security idp status

show security idp memory

- Select **Monitor>Security>IPS>Status** in the J-Web user interface.

[Table 92 on page 472](#) summarizes key output fields in the IDP display.

Table 92: Summary of IDP Status Output Fields

Field	Values	Additional Information
IDP Status		
Status of IDP	Displays the status of the current IDP policy.	—
Up Since	Displays the time from when the IDP policy first began running on the system.	—
Packets/Second	Displays the number of packets received and returned per second.	—
Peak	Displays the maximum number of packets received per second and the time when the maximum was reached.	—
Kbits/Second	Displays the aggregated throughput (kilobits per second) for the system.	—
Peak Kbits	Displays the maximum kilobits per second and the time when the maximum was reached.	—
Latency (Microseconds)	Displays the delay, in microseconds, for a packet to receive and return by a node .	—
Current Policy	Displays the name of the current installed IDP policy.	—
IDP Memory Status		
IDP Memory Statistics	Displays the status of all IDP data plane memory.	—

Table 92: Summary of IDP Status Output Fields (*continued*)

Field	Values	Additional Information
PIC Name	Displays the name of the PIC.	—
Total IDP Data Plane Memory (MB)	Displays the total memory space, in megabytes, allocated for the IDP data plane.	—
Used (MB)	Displays the used memory space, in megabytes, for the data plane.	—
Available (MB)	Displays the available memory space, in megabytes, for the data plane.	—

See Also • *Monitoring Flow Session Statistics*

Monitoring Flow Gate Information

Purpose View information about temporary openings known as pinholes or gates in the security firewall.

Action Select **Monitor>Security>Flow Gate** in the J-Web user interface, or enter the **show security flow gate** command.

Table 93 on page 473 summarizes key output fields in the flow gate display.

Table 93: Summary of Key Flow Gate Output Fields

Field	Values	Additional Information
Flow Gate Information		
Hole	Range of flows permitted by the pinhole.	—
Translated	Tuples used to create the session if it matches the pinhole: <ul style="list-style-type: none"> • Source address and port • Destination address and port 	—
Protocol	Application protocol, such as UDP or TCP.	—
Application	Name of the application.	—
Age	Idle timeout for the pinhole.	—
Flags	Internal debug flags for pinhole.	—
Zone	Incoming zone.	—

Table 93: Summary of Key Flow Gate Output Fields (*continued*)

Field	Values	Additional Information
Reference count	Number of resource manager references to the pinhole.	—
Resource	Resource manager information about the pinhole.	—

- See Also**
- *Monitoring Flow Session Statistics*
 - *Monitoring Firewall Authentication*

Monitoring Firewall Authentication Table

Purpose View information about the authentication table, which divides firewall authentication user information into multiple parts.

Action Select **Monitor>Security>Firewall Authentication>Authentication Table** in the J-Web user interface. To view detailed information about the user with a particular identifier, select the ID on the Authentication Table page. To view detailed information about the user at a particular source IP address, select the Source IP on the Authentication Table page.

Alternatively, enter the following CLI **show** commands:

- **show security firewall-authentication users**
- **show security firewall-authentication users address *ip-address***
- **show security firewall-authentication users identifier *identifier***

[Table 94 on page 474](#) summarizes key output fields in firewall authentication table display.

Table 94: Summary of Key Firewall Authentication Table Output Fields

Field	Values	Additional Information
Firewall authentication users		
Total users in table	Number of users in the authentication table.	—
Authentication table		
ID	Authentication identification number.	—
Source Ip	IP address of the authentication source.	—
Age	Idle timeout for the user.	—
Status	Status of authentication (success or failure).	—
user	Name of the user.	—

Table 94: Summary of Key Firewall Authentication Table Output Fields (*continued*)

Field	Values	Additional Information
Detailed report per ID selected: <i>ID</i>		
Source Zone	Name of the source zone.	—
Destination Zone	Name of the destination zone.	—
profile	Name of the profile.	Users information.
Authentication method	Path chosen for authentication.	—
Policy Id	Policy Identifier.	—
Interface name	Name of the interface.	—
Bytes sent by this user	Number of packets in bytes sent by this user.	—
Bytes received by this user	Number of packets in bytes received by this user.	—
Client-groups	Name of the client group.	—
Detailed report per Source Ip selected		
Entries from Source IP	IP address of the authentication source.	—
Source Zone	Name of the source zone.	—
Destination Zone	Name of the destination zone.	—
profile	Name of the profile.	—
Age	Idle timeout for the user.	—
Status	Status of authentication (success or failure).	—
user	Name of the user.	—
Authentication method	Path chosen for authentication.	—
Policy Id	Policy Identifier.	—
Interface name	Name of the interface.	—
Bytes sent by this user	Number of packets in bytes sent by this user.	—
Bytes received by this user	Number of packets in bytes received by this user.	—
Client-groups	Name of the client group.	—

Monitoring Firewall Authentication History

Purpose View information about the authentication history, which is divided into multiple parts.

Action Select **Monitor>Security>Firewall Authentication>Authentication History** in the J-Web user interface. To view the detailed history of the authentication with this identifier, select the ID on the Firewall Authentication History page. To view a detailed authentication history of this source IP address, select the Source IP on the Firewall Authentication History page.

Alternatively, enter the following CLI **show** commands:

- **show security firewall-authentication history**
- **show security firewall-authentication history address *ip-address***
- **show security firewall-authentication history identifier *identifier***

Table 95 on page 476 summarizes key output fields in firewall authentication history display.

Table 95: Summary of Key Firewall Authentication History Output Fields

Field	Values	Additional Information
History of Firewall Authentication Data		
Total authentications	Number of authentication.	—
History Table		
ID	Identification number.	—
Source Ip	IP address of the authentication source.	—
Start Date	Authentication date.	—
Start Time	Authentication time.	—
Duration	Authentication duration.	—
Status	Status of authentication (success or failure).	—
User	Name of the user.	—
Detail history of selected Id: <i>ID</i>		
Authentication method	Path chosen for authentication.	—
Policy Id	Security policy identifier.	—
Source zone	Name of the source zone.	—

Table 95: Summary of Key Firewall Authentication History Output Fields (*continued*)

Field	Values	Additional Information
Destination Zone	Name of the destination zone.	–
Interface name	Name of the interface.	–
Bytes sent by this user	Number of packets in bytes sent by this user.	–
Bytes received by this user	Number of packets in bytes received by this user.	–
Client-groups	Name of the client group.	–
Detail history of selected Source Ip:Source Ip		
User	Name of the user.	–
Start Date	Authentication date.	–
Start Time	Authentication time.	–
Duration	Authentication duration.	–
Status	Status of authentication (success or failure).	–
Profile	Name of the profile.	–
Authentication method	Path chosen for authentication.	–
Policy Id	Security policy identifier.	–
Source zone	Name of the source zone.	–
Destination Zone	Name of the destination zone.	–
Interface name	Name of the interface.	–
Bytes sent by this user	Number of packets in bytes sent by this user.	–
Bytes received by this user	Number of packets in bytes received by this user.	–
Client-groups	Name of the client group.	–

Monitoring 802.1x

Purpose View information about 802.1X properties.

Action Select **Monitor>Security>802.1x** in the J-Web user interface, or enter the following CLI commands:

- **show dot1x interfaces *interface-name***
- **show dot1x authentication-failed-users**

[Table 96 on page 478](#) summarizes the Dot1X output fields.

Table 96: Summary of Dot1X Output Fields

Field	Values	Additional Information
Select Port	List of ports for selection.	–
Number of connected hosts	Total number of hosts connected to the port.	–
Number of authentication bypassed hosts	Total number of authentication-bypassed hosts with respect to the port.	–
Authenticated Users Summary		
MAC Address	MAC address of the connected host.	–
User Name	Name of the user.	–
Status	Information about the host connection status.	–
Authentication Due	Information about host authentication.	–
Authentication Failed Users Summary		
MAC Address	MAC address of the authentication-failed host.	–
User Name	Name of the authentication-failed user.	–

See Also • [Monitoring Application Firewalls](#)

Related Documentation • [Monitoring Overview on page 7](#)
 • [Monitoring Interfaces on page 426](#)

CHAPTER 27

Monitoring Events, Services and System

- [Monitoring DHCP Client Bindings on page 479](#)
- [Monitoring Events on page 480](#)
- [Monitoring the System on page 482](#)

Monitoring DHCP Client Bindings

Supported Platforms [SRX Series, vSRX](#)

Purpose View information about DHCP client bindings.

Action Select **Monitor>Services>DHCP>Binding** in the J-Web user interface, or enter the **show system services dhcp binding** command.

[Table 97 on page 479](#) summarizes the key output fields in the DHCP client binding displays.

Table 97: Summary of Key DHCP Client Binding Output Fields

Field	Values	Additional Information
IP Address	List of IP addresses the DHCP server has assigned to clients.	—
Hardware Address	Corresponding media access control (MAC) address of the client.	—
Type	Type of binding assigned to the client: dynamic or static.	—
Lease Expires at	Date and time the lease expires, or never for leases that do not expire.	—

Related Documentation

- [Monitoring PPPoE on page 434](#)
- [Understanding DHCP Client Operation](#)

Monitoring Events

Supported Platforms SRX Series, vSRX

Purpose Use the monitoring functionality to view the events page.

Action To monitor events select **Monitor>Events and Alarms>View Events** in the J-Web user interface.



NOTE: When you use an HTTPS connection in the Microsoft Internet Explorer browser to save a report from this page in the J-Web interface, the error message "Internet Explorer was not able to open the Internet site" is displayed. This problem occurs because the Cache-Control: no cache HTTP header is added on the server side and Internet Explorer does not allow you to download the encrypted file with the Cache-Control: no cache HTTP header set in the response from the server.

As a workaround, refer to Microsoft Knowledge Base article 323308, which is available at this URL: <http://support.microsoft.com/kb/323308>. Also, you can alternatively use HTTP in the Internet Explorer browser or use HTTPS in the Mozilla Firefox browser to save a file from this page.

Meaning Table 98 on page 480 summarizes key output fields in the events page.

Table 98: Events Monitoring Page

Field	Value	Additional Information
Events Filter		
System Log File	Specifies the name of the system log file that records errors and events.	-
Process	Specifies the system processes that generate the events to display.	-
Include archived files	Specifies to enable the option to include archived files.	Select to enable.
Date From	Specifies the beginning date range to monitor. Set the date using the calendar pick tool.	-
To	Specifies the end of the date range to monitor. Set the date using the calendar pick tool.	-

Table 98: Events Monitoring Page (*continued*)

Field	Value	Additional Information
Event ID	Specifies the specific ID of the error or event to monitor.	-
Description	Enter a description for the errors or events.	-
Search	Fetches the errors and events specified in the search criteria.	-
Reset	Clears the cache of errors and events that were previously selected.	-
Generate Report	Creates an HTML report based on the specified parameters.	-
Events Detail		
Process	Displays the system process that generated the error or event.	-
Severity	<p>Displays the severity level that indicates how seriously the triggering event affects routing platform functions. Only messages from the facility that are rated at that level or higher are logged. Possible severities and their corresponding color code are:</p> <ul style="list-style-type: none"> • Debug/Info/Notice (Green)—Indicates conditions that are not errors but are of interest or might warrant special handling. • Warning (Yellow) — Indicates conditions that warrant monitoring. • Error (Blue) — Indicates standard error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels. • Critical (Pink) — Indicates critical conditions, such as hard drive errors. • Alert (Orange) — Indicates conditions that require immediate correction, such as a corrupted system database. • Emergency (Red) — Indicates system panic or other conditions that cause the routing platform to stop functioning. 	-
Event ID	Displays the unique ID of the error or event. The prefix on each code identifies the generating software process. The rest of the code indicates the specific event or error.	-
Event Description	Displays a more detailed explanation of the message.	-

Table 98: Events Monitoring Page (*continued*)

Field	Value	Additional Information
Time	Time that the error or event occurred.	-

- Related Documentation**
- [Monitoring Alarms on page 339](#)
 - [Monitoring Security Events by Policy on page 462](#)

Monitoring the System

Supported Platforms [SRX Series, vSRX](#)

The J-Web user interface lets you monitor a device's physical characteristics, current processing status and alarms, and ongoing resource utilization to quickly assess the condition of a device at any time.

On SRX Series devices, the **Dashboard** lets you customize your view by selecting which informational panes to include on the Dashboard.

This section contains the following topics:

- [Monitoring System Properties for SRX Series Devices on page 482](#)
- [Monitoring Chassis Information on page 485](#)
- [System Health Management for SRX Series Devices on page 486](#)

Monitoring System Properties for SRX Series Devices

Purpose View system properties and customize the Dashboard.

When you start the J-Web user interface on an SRX Series device, the interface opens to the Dashboard. At the top and bottom of the page, the Dashboard displays an interactive representation of your device and a current log messages pane. By default, the center panes of the Dashboard display System Information, Resource Utilization, Security Resources, and System Alarms. However, you can customize the Dashboard panes to provide the best overview of your system.

Action To control the content and appearance of the Dashboard:

1. Click the **Preferences** icon at the top-right corner of the page. The Dashboard Preference dialog box appears.
2. Select the types of information you want to display.
3. (Optional) Specify the Automatically Refresh Data option to specify how often you want the data on the Dashboard to be refreshed.

4. Click **OK** to save the configuration or **Cancel** to clear it.
5. On the Dashboard, minimize, maximize, or drag the individual information panes to customize the display as needed.

Chassis View—Displays an image of the device chassis, including line cards, link states, errors, individual PICs, FPCs, fans, and power supplies.

You can use the Chassis View to link to corresponding configuration and monitoring pages for the device. To link to interface configuration pages for a selected port from the Chassis View, right-click the port in the device image and choose one of the following options:

- Chassis Information—Links to the Chassis page.
- Configure Port: *Port-name*—Links to the interfaces configuration page for the selected port.
- Monitor Port: *Port-name*—Links to the monitor interfaces page for the selected port.

System Identification—Displays the device's serial number, hostname, current software version, the BIOS version, the amount of time since the device was last booted, and the system's time.



NOTE:

- To view the BIOS version under system identification, delete your browser cookies.
- The hostname that appears in this pane is defined using the `set system hostname` command.

On SRX Series devices, security logs were always timestamped using the UTC time zone by running `set system time-zone utc` and `set security log utc-timestamp` CLI commands. Now, time zone can be defined using the local time zone by running the `set system time-zone time-zone` command to specify the local time zone that the system should use when timestamping the security logs.

Resource Utilization—Provides a graphic representation of resource use. Each bar represents the percentage of CPU, memory, or storage utilization for the data plane or the control plane.

Security Resources—Provides the maximum, configured, and active sessions; firewall and VPN policies; and IPsec VPNs. Click **Sessions**, **FW/VPN Policies**, or **IPsec VPNs** for detailed statistics about each category.

System Alarms—Indicates a missing rescue configuration or software license, where valid. System alarms are preset and cannot be modified.

File Usage—Displays the usage statistics for log files, temporary files, crash (core) files, and database files.

Login Sessions—Provides a list of all currently logged in sessions. The display includes user credentials, login time, and idle time for each session.

Chassis Status—Provides a snapshot of the current physical condition of the device, including temperature and fan status.

Storage Usage—Displays the storage usage report in detail.

Threat Activity—Provides information about the most current threats received on the device.

Message Logs—Displays log messages and errors. You can clear old logs from the Message Logs pane by clicking the Clear button.

To control the information that is displayed in the Chassis View, use the following options:

- To view an image of the front of the device, right-click the image and choose **View Front**.
- To view an image of the back of the device, right-click the image and choose **View Rear**.
- To enlarge or shrink the device view, use the **Zoom** bar.
- To return the device image to its original position and size, click **Reset**.



NOTE: To use the Chassis View, a recent version of Adobe Flash that supports ActionScript and AJAX (Version 9) must be installed. Also note that the Chassis View appears by default on the Dashboard page. You can enable or disable it using options in the Dashboard Preference dialog box. Clearing cookies in Internet Explorer also causes the Chassis View appear on the Dashboard page.

To return to the Dashboard at any time, select **Dashboard** in the J-Web user interface.

Alternatively, you can view system properties by entering the following **show** commands in the CLI:

- **show system uptime**
- **show system users**
- **show system storage**
- **show version**
- **show chassis hardware**

Monitoring Chassis Information

Purpose View chassis properties, which include the status of hardware components on the device.

Action To view these chassis properties, select **Monitor>System View>Chassis Information** in the J-Web user interface.



CAUTION: Do not install a combination of Physical Interface Modules (PIMs) in a single chassis that exceeds the maximum power and heat capacity of the chassis. If power management is enabled, PIMs that exceed the maximum power and heat limits remain offline when the chassis is powered on. To check PIM power and heat status, use the `show chassis fpc` and `show chassis power-ratings` commands.

The Chassis Information page displays the following types of information:

- Routing Engine Details—This section of the page includes the following tabs:
 - Master—Master tab displays information about the routing engine, including the routing engine module, model number, version, part number, serial number, memory utilization, temperature, and start time. Additionally, this tab displays the CPU load averages for the last 1, 5, and 15 minutes.
 - Backup—If a backup routing engine is available, the Backup tab displays the routing engine module, model number, version, part number, serial number, memory utilization, temperature, and start time. Additionally, this tab displays the CPU load averages for the last 1, 5, and 15 minutes.



NOTE: If you need to contact customer support about the device chassis, supply them with the version and serial number displayed in the Routing Engine Details section of the page.

- Power and Fan Tray Details—This Details section of the page includes the following tabs:
 - Power—Power tab displays the names of the device's power supply units and their statuses.
 - Fan—Fan tab displays the names of the device's fans and their speeds (normal or high). (The fan speeds are adjusted automatically according to the current temperature.)
- Chassis Component Details—This section of the page includes the following tabs:

- General—General tab displays the version number, part number, serial number, and description of the selected device component.
- Temperature—Temperature tab displays the temperature of the selected device component (if applicable).
- Resource—Resource tab displays the state, total CPU DRAM, and start time of the selected device component (if applicable).



NOTE: On some devices, you can have an FPC state as “offline.” You might want to put an FPC offline because of an error or if the FPC is not responding. You can put the FPC offline by using the CLI command `request chassis fpc slot number offline`.

- Sub-Component—Sub-Component tab displays information about the device's sub-components (if applicable). Details include the sub-component's version, part number, serial number, and description.

To control which component details appear, select a hardware component from the **Select component** list.

Alternatively, you can view chassis details by entering the following **show** commands in the CLI configuration editor:

- **show chassis hardware**
- **show chassis routing-engine**
- **show chassis environment**
- **show chassis redundant-power-supply**
- **show redundant-power-supply status**

- See Also**
- *Monitoring Cluster Status*
 - *Monitoring Cluster Statistics*

System Health Management for SRX Series Devices

Purpose Tracking the utilization of critical resources in the system ensures that all parameters are within normal limits and the system remains functional.

In the event of a malfunction caused by abnormal resource usage, the system health management feature provides the right diagnostic information to identify the source of the problem.

When the system health management action is configured by the user, the system takes appropriate monitoring, preventive, and recovery actions to ensure that the system is accessible. The system configuration might be updated based on the information collected

by system health management feature to ensure that the system stays in the normal operating environment. For example, when a system runs out of memory, then the configuration associated with applications identified to be consuming memory resources can be updated to bring down the memory resource consumption.

Action The system health management feature periodically monitors critical system resources against configurable thresholds. The resources that can be monitored include CPU usage, memory, storage, open-file-descriptor, process-count, and temperature. The system health management feature collects usage information for each resource at the configured interval and compares it against the three levels of thresholds: moderate, high, and critical. Based on the configurations, appropriate action is taken.

The intervals, thresholds, and action are associated with system health management and can be configured at both the resource level and the global level. Configurable and default levels are as follows:

- Default configuration level— Default configuration is applied when system health monitoring is enabled, and neither a global nor a resource-specific configuration is present.
- Global configuration level—Configuration that is applied to resources when no resource-specific configuration is available.
- Resource-specific configuration level—Configuration that, if available, overrides both the global and the default configurations.

Per-resource configurations take precedence over the global configuration, and a global configuration takes precedence over the defaults.

When resource usage exceeds the configured thresholds, the system collects information that can be used to find the source of the increased usage and saves it in history for analysis and action.

When resource utilization exceeds the high threshold, a minor system alarm is generated, and the alarm LED lights yellow. When resource utilization exceeds the critical threshold, a major alarm is generated, and the alarm LED lights red.

An SNMP trap is also sent to the remote monitoring server (NMS) for all events that exceed the threshold.

To enable the system health monitor, use the **set snmp health-monitor routing engine** command. You can view system properties by using CLI show commands.

Related Documentation

- [Monitoring Overview on page 7](#)
- [Monitoring Interfaces on page 426](#)

Monitoring Unified Threat Management Features

- [Monitoring Antivirus Scan Engine Status on page 489](#)
- [Monitoring Antivirus Scan Results on page 490](#)
- [Monitoring Antivirus Session Status on page 492](#)
- [Monitoring Content Filtering Configurations on page 493](#)
- [Monitoring Reports on page 494](#)
- [Monitoring Web Filtering Configurations on page 501](#)

Monitoring Antivirus Scan Engine Status

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX](#)

Purpose The Monitoring Antivirus Scan Engine Status is not supported from Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1 onwards. For previous releases, using the CLI, you can view the following scan engine status items:

Antivirus license key status

- View license expiration dates.

Scan engine status and settings

- View last action result.
- View default file extension list.

Antivirus pattern update server settings

- View update URL (HTTP or HTTPS-based).
- View update interval.

Antivirus pattern database status

- View auto update status.
- View last result of database loading.

- If the download completes, view database version timestamp virus record number.
- If the download fails, view failure reason.

Action In the CLI, enter the `user@host> show security utm anti-virus status` command.

Example status result:

```
AV Key Expire Date: 03/01/2010 00:00:00
Update Server: http://update.juniper-updates.net/AV/device-name
interval: 60 minutes
auto update status: next update in 12 minutes
last result: new database loaded
AV signature version: 12/21/2008 00:35 GMT, virus records: 154018
Scan Engine Info: last action result: No error(0x00000000)
```

Release History Table

Release	Description
15.1X49-D10	The Monitoring Antivirus Scan Engine Status is not supported from Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1 onwards.

Related Documentation

- *Understanding the Full Antivirus Scan Engine*

Monitoring Antivirus Scan Results

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [SRX345](#), [SRX550M](#), [vSRX](#)

Purpose The Monitoring Antivirus Scan Results are not supported from Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1 onwards. For previous releases, view statistics for antivirus requests, scan results, and fallback counters.

Scan requests provide

- The total number of scan request forwarded to the engine.
- The number of scan request being pre-windowed.
- The number of scan requests using scan-all mode.
- The number of scan requests using scan-by-extension mode.

Scan code counters provide

- Number of clean files.
- Number of infected files.
- Number of password protected files.
- Number of decompress layers.
- Number of corrupt files.

- When the engine is out of resources.
- When there is an internal error.

Fallback applied status provides either a log-and-permit or block result when the following has occurred

- Scan engine not ready.
- Maximum content size reached.
- Too many requests.
- Password protected file found.
- Decompress layer too large.
- Corrupt file found.
- Timeout occurred.
- Out of resources.
- Other.

Action To view antivirus scan results using the CLI editor, enter the **user@host> show security utm anti-virus statistics status** command.

To view antivirus scan results using J-Web:

1. Select **Monitor>Security>UTM>Anti-Virus**.

The following information becomes viewable in the right pane.

Antivirus license key status

- View license expiration dates.

Antivirus pattern update server settings

- View update URL (HTTP or HTTPS-based).
- View update interval.

Antivirus pattern database status

- View auto update status.
- View last result of database loading.
- If the download completes, view database version timestamp virus record number.
- If the download fails, view failure reason.

Antivirus statistics provide

- The number of scan request being pre-windowed.
- The total number of scan request forwarded to the engine.

- The number of scan requests using scan-all mode.
- The number of scan requests using scan-by-extension mode.

Scan code counters provide

- Number of clean files.
- Number of infected files.
- Number of password protected files.
- Number of decompress layers.
- Number of corrupt files.
- When the engine is out of resources.
- When there is an internal error.

Fallback applied status provides either a log-and-permit or block result when the following has occurred

- Scan engine not ready.
- Password protected file found.
- Decompress layer too large.
- Corrupt file found.
- Out of resources.
- Timeout occurred.
- Maximum content size reached.
- Too many requests.
- Other.

2. You can click the **Clear Anti-Virus Statistics** button to clear all current viewable statistics and begin collecting new statistics.

Release History Table

Release	Description
15.1X49-D10	The Monitoring Antivirus Scan Results are not supported from Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1 onwards.

Monitoring Antivirus Session Status

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX

Purpose The Monitoring Antivirus Session Status is not supported from Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1 onwards. For previous releases, using the CLI, you can view the following session status items:

Antivirus session status displays a snapshot of current antivirus sessions. It includes

- Maximum supported antivirus session numbers.
- Total allocated antivirus session numbers.
- Total freed antivirus session numbers.
- Current active antivirus session numbers.

Action In the CLI, enter the `user@host> show security utm session status` command.

Release History Table

Release	Description
15.1X49-D10	The Monitoring Antivirus Session Status is not supported from Junos OS Release 15.1X49-D10 and Junos OS Release 17.3R1 onwards.

Monitoring Content Filtering Configurations

Supported Platforms [SRX Series, vSRX](#)

Purpose View content filtering statistics.

Action To view content filtering statistics in the CLI, enter the `user@host > show security utm content-filtering statistics` command.

The content filtering **show statistics** command displays the following information:

```
Base on command list: # Blocked
Base on mime list: # Blocked
Base on extension list: # Blocked
ActiveX plugin: # Blocked
Java applet: # Blocked
EXE files: # Blocked
ZIP files: # Blocked
HTTP cookie: # Blocked
```

To view content filtering statistics using J-Web:

1. Select **Clear Content filtering statisticsMonitor>Security>UTM>Content FilteringMonitor>Security>UTM>Content Filtering**.

The following statistics become viewable in the right pane.

```
Base on command list: # Passed # Blocked
Base on mime list: # Passed # Blocked
Base on extension list: # Passed # Blocked
ActiveX plugin: # Passed # Blocked
Java applet: # Passed # Blocked
EXE files: # Passed # Blocked
ZIP files: # Passed # Blocked
HTTP cookie: # Passed # Blocked
```

-
2. You can click **Clear Content filtering statistics** to clear all current viewable statistics and begin collecting new statistics.

Monitoring Reports

Supported Platforms [SRX Series, vSRX](#)

On-box reporting offers a comprehensive reporting facility where your security management team can spot a security event when it occurs, immediately access and review pertinent details about the event, and quickly decide appropriate remedial action. The J-Web reporting feature provides one- or two-page reports that are equivalent to a compilation of numerous log entries.

This section contains the following topics:

- [Threats Monitoring Report on page 494](#)
- [Traffic Monitoring Report on page 499](#)

Threats Monitoring Report

Purpose Use the Threats Report to monitor general statistics and activity reports of current threats to the network. You can analyze logging data for threat type, source and destination details, and threat frequency information. The report calculates, displays, and refreshes the statistics, providing graphic presentations of the current state of the network.

Action To view the Threats Report:

1. Click **Threats Report** in the bottom right of the Dashboard, or select **Monitor>Reports>Threats** in the J-Web user interface. The Threats Report appears.
2. Select one of the following tabs:
 - **Statistics** tab. See [Table 99 on page 494](#) for a description of the page content.
 - **Activities** tab. See [Table 100 on page 497](#) for a description of the page content.

Table 99: Statistics Tab Output in the Threats Report

Field	Description
-------	-------------

General Statistics Pane

Table 99: Statistics Tab Output in the Threats Report (*continued*)

Field	Description
Threat Category	One of the following categories of threats: <ul style="list-style-type: none"> • Traffic • IDP • Content Security <ul style="list-style-type: none"> • Antivirus • Antispam • Web Filter—Click the Web filter category to display counters for 39 subcategories. • Content Filter • Firewall Event
Severity	Severity level of the threat: <ul style="list-style-type: none"> • emerg • alert • crit • err • warning • notice • info • debug
Hits in past 24 hours	Number of threats encountered per category in the past 24 hours.
Hits in current hour	Number of threats encountered per category in the last hour.
Threat Counts in the Past 24 Hours	
By Severity	Graph representing the number of threats received each hour for the past 24 hours sorted by severity level.
By Category	Graph representing the number of threats received each hour for the past 24 hours sorted by category.
X Axis	Twenty-four hour span with the current hour occupying the right-most column of the display. The graph shifts to the left every hour.
Y Axis	Number of threats encountered. The axis automatically scales based on the number of threats encountered.
Most Recent Threats	
Threat Name	Names of the most recent threats. Depending on the threat category, you can click the threat name to go to a scan engine site for a threat description.

Table 99: Statistics Tab Output in the Threats Report (*continued*)

Field	Description
Category	Category of each threat: <ul style="list-style-type: none"> Traffic IDP Content Security <ul style="list-style-type: none"> Antivirus Antispam Web Filter Content Filter Firewall Event
Source IP/Port	Source IP address (and port number, if applicable) of the threat.
Destination IP/Port	Destination IP address (and port number, if applicable) of the threat.
Protocol	Protocol name of the threat.
Description	Threat identification based on the category type: <ul style="list-style-type: none"> Antivirus—URL Web filter—category Content filter—reason Antispam—sender e-mail
Action	Action taken in response to the threat.
Hit Time	Time the threat occurred.
Threat Trend in past 24 hours	
Category	Pie chart graphic representing comparative threat counts by category: <ul style="list-style-type: none"> Traffic IDP Content Security <ul style="list-style-type: none"> Antivirus Antispam Web Filter Content Filter Firewall Event
Web Filter Counters Summary	
Category	Web filter count broken down by up to 39 subcategories. Clicking on the Web filter listing in the General Statistics pane opens the Web Filter Counters Summary pane.
Hits in past 24 hours	Number of threats per subcategory in the last 24 hours.

Table 99: Statistics Tab Output in the Threats Report (*continued*)

Field	Description
Hits in current hour	Number of threats per subcategory in the last hour.

Table 100: Activities Tab Output in the Threats Report

Field	Function
Most Recent Virus Hits	
Threat Name	Name of the virus threat. Viruses can be based on services, like Web, FTP, or e-mail, or based on severity level.
Severity	Severity level of each threat: <ul style="list-style-type: none"> • emerg • alert • crit • err • warning • notice • info • debug
Source IP/Port	IP address (and port number, if applicable) of the source of the threat.
Destination IP/Port	IP address (and port number, if applicable) of the destination of the threat.
Protocol	Protocol name of the threat.
Description	Threat identification based on the category type: <ul style="list-style-type: none"> • Antivirus—URL • Web filter—category • Content filter—reason • Antispam—sender e-mail
Action	Action taken in response to the threat.
Last Hit Time	Last time the threat occurred.
Most Recent Spam E-Mail Senders	
From e-mail	E-mail address that was the source of the spam.

Table 100: Activities Tab Output in the Threats Report (*continued*)

Field	Function
Severity	Severity level of the threat: <ul style="list-style-type: none"> • emerg • alert • crit • err • warning • notice • info • debug
Source IP	IP address of the source of the threat.
Action	Action taken in response to the threat.
Last Send Time	Last time that the spam e-mail was sent.
Recently Blocked URL Requests	
URL	URL request that was blocked.
Source IP/Port	IP address (and port number, if applicable) of the source.
Destination IP/Port	IP address (and port number, if applicable) of the destination.
Hits in current hour	Number of threats encountered in the last hour.
Most Recent IDP Attacks	
Attack	
Severity	Severity of each threat: <ul style="list-style-type: none"> • emerg • alert • crit • err • warning • notice • info • debug
Source IP/Port	IP address (and port number, if applicable) of the source.
Destination IP/Port	IP address (and port number, if applicable) of the destination.
Protocol	Protocol name of the threat.

Table 100: Activities Tab Output in the Threats Report (*continued*)

Field	Function
Action	Action taken in response to the threat.
Last Send Time	Last time the IDP threat was sent.

- See Also**
- [Traffic Monitoring Report on page 499](#)
 - [Monitoring Address Pools on page 423](#)

Traffic Monitoring Report

Purpose Monitor network traffic by reviewing reports of flow sessions over the past 24 hours. You can analyze logging data for connection statistics and session usage by a transport protocol.

Action To view network traffic in the past 24 hours, select **Monitor>Reports>Traffic** in the J-Web user interface. See [Table 101 on page 499](#) for a description of the report.

Table 101: Traffic Report Output

Field	Description
Sessions in Past 24 Hours per Protocol	
Protocol Name	Name of the protocol. To see hourly activity by protocol, click the protocol name and review the “Protocol activities chart” in the lower pane. <ul style="list-style-type: none"> • TCP • UDP • ICMP
Total Session	Total number of sessions for the protocol in the past 24 hours.
Bytes In (KB)	Total number of incoming bytes in KB.
Bytes Out (KB)	Total number of outgoing bytes in KB.
Packets In	Total number of incoming packets.
Packets Out	Total number of outgoing packets.
Most Recently Closed Sessions	
Source IP/Port	Source IP address (and port number, if applicable) of the closed session.
Destination IP/Port	Destination IP address (and port number, if applicable) of the closed session.

Table 101: Traffic Report Output (*continued*)

Field	Description
Protocol	Protocol of the closed session. <ul style="list-style-type: none"> • TCP • UDP • ICMP
Bytes In (KB)	Total number of incoming bytes in KB.
Bytes Out (KB)	Total number of outgoing bytes in KB.
Packets In	Total number of incoming packets.
Packets Out	Total number of outgoing packets.
Timestamp	The time the session was closed.
Protocol Activities Chart	
Bytes In/Out	Graphic representation of traffic as incoming and outgoing bytes per hour. The byte count is for the protocol selected in the Sessions in Past 24 Hours per Protocol pane. Changing the selection causes this chart to refresh immediately.
Packets In/Out	Graphic representation of traffic as incoming and outgoing packets per hour. The packet count is for the protocol selected in the Sessions in Past 24 Hours per Protocol pane. Changing the selection causes this chart to refresh immediately.
Sessions	Graphic representation of traffic as the number of sessions per hour. The session count is for the protocol selected in the Sessions in Past 24 Hours per Protocol pane. Changing the selection causes this chart to refresh immediately.
X Axis	One hour per column for 24 hours.
Y Axis	Byte, packet, or session count.
Protocol Session Chart	
Sessions by Protocol	Graphic representation of the traffic as the current session count per protocol. The protocols displayed are TCP, UDP, and ICMP.

See Also • [Threats Monitoring Report on page 494](#)

Related Documentation • [Monitoring Overview on page 7](#)
• [Monitoring Interfaces on page 426](#)

Monitoring Web Filtering Configurations

Supported Platforms [SRX Series, vSRX](#)

Purpose View Web-filtering statistics.

Action To view Web-filtering statistics using the CLI, enter the following commands:

```
user@host> show security utm web-filtering status
user@host> show security utm web-filtering statistics
```

To view Web-filtering statistics using J-Web:

1. Select **Clear Web Filtering Statistics**.

The following information is displayed in the right pane.

```
Total Requests: #
White List Hit: #
Black List Hit: #
Queries to Server: #
Server Reply Permit: #
Server Reply Block: #
Custom Category Permit: #
Custom Category Block: #
Cache Hit Permit: #
Cache Hit Block: #
Web Filtering Session Total: #
Web Filtering Session Inuse: #
Fall Back: Log-and-Permit Block
Default # #
Timeout # #
Server-Connectivity # #
Too-Many-Requests # #
```

2. You can click the **Clear Web Filtering Statistics** button to clear all current viewable statistics and begin collecting new statistics.

Related Documentation

- [Web Filtering Overview](#)
- [Example: Configuring Enhanced Web Filtering](#)

Monitoring VPNs

- [Monitoring VPNs on page 503](#)

Monitoring VPNs

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring IKE Gateway Information on page 503](#)
- [Monitoring IPsec VPN—Phase I on page 507](#)
- [Monitoring IPsec VPN—Phase II on page 508](#)
- [Monitoring IPsec VPN Information on page 509](#)

Monitoring IKE Gateway Information

Purpose View information about IKE security associations (SAs).

Action Select **Monitor>IPSec VPN>IKE Gateway** in the J-Web user interface. To view detailed information for a particular SA, select the IKE SA index on the IKE gateway page.

Alternatively, enter the following CLI commands:

- **show security ike security-associations**
- **show security ike security-associations index *index-id* detail**

[Table 102 on page 503](#) summarizes key output fields in the IKE gateway display.

Table 102: Summary of Key IKE SA Information Output Fields

Field	Values	Additional Information
IKE Security Associations		
IKE SA Index	Index number of an SA.	This number is an internally generated number you can use to display information about a single SA.
Remote Address	IP address of the destination peer with which the local peer communicates.	–

Table 102: Summary of Key IKE SA Information Output Fields (*continued*)

Field	Values	Additional Information
State	State of the IKE security associations: <ul style="list-style-type: none"> • DOWN—SA has not been negotiated with the peer. • UP—SA has been negotiated with the peer. 	—
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.	—
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received.	A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Mode	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are: <ul style="list-style-type: none"> • Main—The exchange is done with six messages. This mode, or exchange type, encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • Aggressive—The exchange is done with three messages. This mode, or exchange type, does not encrypt the payload, leaving the identity of the neighbor unprotected. 	—
IKE Security Association (SA) Index		
IKE Peer	IP address of the destination peer with which the local peer communicates.	—
IKE SA Index	Index number of an SA.	This number is an internally generated number you can use to display information about a single SA.
Role	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.	—
State	State of the IKE security associations: <ul style="list-style-type: none"> • DOWN—SA has not been negotiated with the peer. • UP—SA has been negotiated with the peer. 	—

Table 102: Summary of Key IKE SA Information Output Fields (*continued*)

Field	Values	Additional Information
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.	–
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received.	A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Exchange Type	<p>Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between themselves. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are:</p> <ul style="list-style-type: none"> • Main—The exchange is done with six messages. This mode, or exchange type, encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • Aggressive—The exchange is done with three messages. This mode, or exchange type, does not encrypt the payload, leaving the identity of the neighbor unprotected. 	–
Authentication Method	Path chosen for authentication.	–
Local	Address of the local peer.	–
Remote	Address of the remote peer.	–
Lifetime	Number of seconds remaining until the IKE SA expires.	–

Table 102: Summary of Key IKE SA Information Output Fields (*continued*)

Field	Values	Additional Information
Algorithm	<p>IKE algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:</p> <ul style="list-style-type: none"> • Authentication—Type of authentication algorithm used. <ul style="list-style-type: none"> • sha1—Secure Hash Algorithm 1 (SHA-1) authentication. • md5—MD5 authentication. • Encryption—Type of encryption algorithm used. <ul style="list-style-type: none"> • aes-256-cbc—Advanced Encryption Standard (AES) 256-bit encryption. • aes-192-cbc—Advanced Encryption Standard (AES) 192-bit encryption. • aes-128-cbc—Advanced Encryption Standard (AES) 128-bit encryption. • 3des-cbc—3 Data Encryption Standard (DES) encryption. • des-cbc—Data Encryption Standard (DES) encryption. • Pseudorandom function—Cryptographically secure pseudorandom function family. 	—
Traffic Statistics	<p>Traffic statistics include the following:</p> <ul style="list-style-type: none"> • Input bytes—The number of bytes presented for processing by the device. • Output bytes—The number of bytes actually processed by the device. • Input packets—The number of packets presented for processing by the device. • Output packets—The number of packets actually processed by the device. 	—
IPsec security associations	<ul style="list-style-type: none"> • number created—The number of SAs created. • number deleted—The number of SAs deleted. 	—
Role	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.	—
Message ID	Message identifier.	—
Local identity	Specifies the identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as any of the following: IPv4 address, fully qualified domain name, e-mail address, or distinguished name.	—

Table 102: Summary of Key IKE SA Information Output Fields (*continued*)

Field	Values	Additional Information
Remote identity	IPv4 address of the destination peer gateway.	—

Monitoring IPsec VPN—Phase I

Purpose View IPsec VPN Phase I information.

Action Select **Monitor>IPSec VPN>Phase I** in the J-Web user interface.

Table 103 on page 507 describes the available options for monitoring IPsec VPN-Phase I.

Table 103: IPsec VPN—Phase I Monitoring Page

Field	Values	Additional Information
IKE SA Tab Options		
IKE Security Associations		
SA Index	Index number of an SA.	—
Remote Address	IP address of the destination peer with which the local peer communicates.	—
State	State of the IKE security associations: <ul style="list-style-type: none"> DOWN—SA has not been negotiated with the peer. UP—SA has been negotiated with the peer. 	—
Initiator Cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.	—
Responder Cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received.	A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.

Table 103: IPsec VPN—Phase I Monitoring Page (*continued*)

Field	Values	Additional Information
Mode	<p>Negotiation method agreed upon by the two IPsec endpoints, or peers, used to exchange information. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are:</p> <ul style="list-style-type: none"> • Main—The exchange is done with six messages. This mode, or exchange type, encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • Aggressive—The exchange is done with three messages. This mode, or exchange type, does not encrypt the payload, leaving the identity of the neighbor unprotected. 	—

See Also • [Monitoring IPsec VPN—Phase II on page 508](#)

Monitoring IPsec VPN—Phase II

Purpose View IPsec VPN Phase II information.

Action Select **Monitor>IPSec VPN>Phase II** in the J-Web user interface.

[Table 104 on page 508](#) describes the available options for monitoring IPsec VPN-Phase II.

Table 104: IPsec VPN—Phase II Monitoring Page

Field	Values	Additional Information
Statistics Tab Details		
By bytes	Provides total number of bytes encrypted and decrypted by the local system across the IPsec tunnel.	—
By packets	Provides total number of packets encrypted and decrypted by the local system across the IPsec tunnel.	—
IPsec Statistics	Provides details of the IPsec statistics.	—
IPsec SA Tab Details		
IPsec Security Associations		

Table 104: IPsec VPN—Phase II Monitoring Page (*continued*)

Field	Values	Additional Information
ID	Index number of the SA.	—
Gateway/Port	IP address of the remote gateway/port.	—
Algorithm	<p>Cryptography scheme used to secure exchanges between peers during the IKE Phase II negotiations:</p> <ul style="list-style-type: none"> An authentication algorithm used to authenticate exchanges between the peers. Options are hmac-md5-95 or hmac-sha1-96. 	—
SPI	<p>Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase I and Phase II.</p>	—
Life	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.	—
Monitoring	Specifies if VPN-Liveliness Monitoring has been enabled/disabled. Enabled - 'U ', Disabled- '—'	—
Vsys	Specifies the root system.	—

See Also • [Monitoring IPsec VPN—Phase I on page 507](#)

Monitoring IPsec VPN Information

Purpose View information about IPsec security (SAs).

Action Select **Monitor>IPSec VPN>IPsec VPN** in the J-Web user interface. To view the IPsec statistics information for a particular SA, select the IPsec SA ID value on the IPsec VPN page.

Alternatively, enter the following CLI commands:

- `show security ipsec security-associations`
- `show security ipsec statistics`

Table 105 on page 510 summarizes key output fields in the IPsec VPN display.

Table 105: Summary of Key IPsec VPN Information Output Fields

Field	Values	Additional Information
IPsec Security Associations		
Total configured SA	Total number of IPsec security associations (SAs) configured on the device.	—
ID	Index number of the SA.	—
Gateway	IP address of the remote gateway.	—
Port	If Network Address Translation (NAT-T) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.	—
Algorithm	<p>Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations:</p> <ul style="list-style-type: none"> An authentication algorithm used to authenticate exchanges between the peers. Options are hmac-md5-95 or hmac-sha1-96. An encryption algorithm used to encrypt data traffic. Options are 3des-cbc, aes-128-cbc, aes-192-cbc, aes-256-cbc, or des-cbc. 	—
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.	—
Life: sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.	—
State	<p>State has two options, Installed and Not Installed.</p> <ul style="list-style-type: none"> Installed—The security association is installed in the security association database. Not Installed—The security association is not installed in the security association database. 	For transport mode, the value of State is always Installed .
Vsys	The root system.	—
IPsec Statistics Information		

Table 105: Summary of Key IPsec VPN Information Output Fields (*continued*)

Field	Values	Additional Information
ESP Statistics	<p>Encapsulation Security Protocol (ESP) statistics include the following:</p> <ul style="list-style-type: none"> • Encrypted bytes—Total number of bytes encrypted by the local system across the IPsec tunnel. • Decrypted bytes—Total number of bytes decrypted by the local system across the IPsec tunnel. • Encrypted packets—Total number of packets encrypted by the local system across the IPsec tunnel. • Decrypted packets—Total number of packets decrypted by the local system across the IPsec tunnel. 	—
AH Statistics	<p>Authentication Header (AH) statistics include the following:</p> <ul style="list-style-type: none"> • Input bytes—The number of bytes presented for processing by the device. • Output bytes—The number of bytes actually processed by the device. • Input packets—The number of packets presented for processing by the device. • Output packets—The number of packets actually processed by the device. 	—
Errors	<p>Errors include the following</p> <ul style="list-style-type: none"> • AH authentication failures—Total number of authentication header (AH) failures. An AH failure occurs when there is a mismatch of the authentication header in a packet transmitted across an IPsec tunnel. • Replay errors—Total number of replay errors. A replay error is generated when a duplicate packet is received within the replay window. • ESP authentication failures—Total number of Encapsulation Security Payload (ESP) failures. An ESP failure occurs when there is an authentication mismatch in ESP packets. • ESP decryption failures—Total number of ESP decryption errors. • Bad headers—Total number of invalid headers detected. • Bad trailers—Total number of invalid trailers detected. 	—
Details for IPsec SA Index: <i>ID</i>		
Virtual System	The root system.	—

Table 105: Summary of Key IPsec VPN Information Output Fields (*continued*)

Field	Values	Additional Information
Local Gateway	Gateway address of the local system.	—
Remote Gateway	Gateway address of the remote system.	—
Local identity	Specifies the identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as any of the following: IPv4 address, fully qualified domain name, e-mail address, or distinguished name.	—
Remote identity	IPv4 address of the destination peer gateway.	—
Df bit	State of the don't fragment bit— set or cleared .	—
Policy name	Name of the applicable policy.	—
Direction	Direction of the security association— inbound , or outbound .	—
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.	—
Mode	Mode of the security association. Mode can be transport or tunnel. <ul style="list-style-type: none"> • transport—Protects host-to-host connections. • tunnel—Protects connections between security gateways. 	—
Type	Type of the security association, either manual or dynamic . <ul style="list-style-type: none"> • manual—Security parameters require no negotiation. They are static and are configured by the user. • dynamic—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode. 	—

Table 105: Summary of Key IPsec VPN Information Output Fields (*continued*)

Field	Values	Additional Information
State	<p>State has two options, Installed, and Not Installed.</p> <ul style="list-style-type: none"> • Installed—The security association is installed in the security association database. • Not Installed—The security association is not installed in the security association database. 	For transport mode, the value of State is always Installed .
Protocol	<p>Protocol supported:</p> <ul style="list-style-type: none"> • Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH). • Tunnel mode supports ESP and AH. <ul style="list-style-type: none"> • Authentication—Type of authentication used. • Encryption—Type of encryption used. 	—
Authentication/Encryption	<ul style="list-style-type: none"> • Authentication—Type of authentication algorithm used. <ul style="list-style-type: none"> • sha1—Secure Hash Algorithm 1 (SHA-1) authentication. • md5—MD5 authentication. • Encryption—Type of encryption algorithm used. <ul style="list-style-type: none"> • aes-256-cbc—Advanced Encryption Standard (AES) 256-bit encryption. • aes-192-cbc—Advanced Encryption Standard (AES) 192-bit encryption. • aes-128-cbc—Advanced Encryption Standard (AES) 128-bit encryption. • 3des-cbc—3 Data Encryption Standard (DES) encryption. • des-cbc—Data Encryption Standard (DES) encryption. 	—
Soft Lifetime	<p>The soft lifetime informs the IPsec key management system that the SA is about to expire.</p> <ul style="list-style-type: none"> • Expires in seconds—Number of seconds left until the SA expires. • Expires in kilobytes—Number of kilobytes left until the SA expires. 	Each lifetime of a security association has two display options, hard and soft , one of which must be present for a dynamic security association. This allows the key management system to negotiate a new SA before the hard lifetime expires.
Hard Lifetime	<p>The hard lifetime specifies the lifetime of the SA.</p> <ul style="list-style-type: none"> • Expires in seconds—Number of seconds left until the SA expires. • Expires in kilobytes—Number of kilobytes left until the SA expires. 	—

Table 105: Summary of Key IPsec VPN Information Output Fields (*continued*)

Field	Values	Additional Information
Anti Replay Service	State of the service that prevents packets from being replayed. It can be Enabled or Disabled .	–
Replay Window Size	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled.	The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.

- Related Documentation**
- [Monitoring Overview on page 7](#)
 - [Monitoring Interfaces on page 426](#)

PART 8

Resource Monitoring of Memory Regions and Types Using CLI and SNMP Queries

- [Effective Troubleshooting of System Performance With Resource Monitoring Methodology on page 517](#)

CHAPTER 30

Effective Troubleshooting of System Performance With Resource Monitoring Methodology

- [Resource Monitoring Usage Computation Overview on page 517](#)
- [Diagnosing and Debugging System Performance by Configuring Memory Resource Usage Monitoring on MX Series Routers on page 520](#)
- [Troubleshooting the Mismatch of jnxNatObjects Values for MS-DPC and MS-MIC on page 523](#)
- [Managed Objects for Ukernel Memory for a Packet Forwarding Engine in an FPC Slot on page 524](#)
- [Managed Objects for Packet Forwarding Engine Memory Statistics Data on page 524](#)
- [Managed Objects for Next-Hop, Jtree, and Firewall Filter Memory for a Packet Forwarding Engine in an FPC Slot on page 525](#)
- [jnxPfeMemoryErrorsTable on page 525](#)
- [pfeMemoryErrors on page 526](#)

Resource Monitoring Usage Computation Overview

Supported Platforms MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

You can configure the resource monitoring capability using both the CLI and SNMP MIB queries. You can employ this utility to provision sufficient headroom (memory space limits that are set for the application or virtual router) for monitoring the health and operating efficiency of DPCs and MPCs. You can also analyze and view the usage or consumption of memory for the jtree memory type and for contiguous pages, double words, and free memory pages. The jtree memory on all MX Series router Packet Forwarding Engines has two segments: one segment primarily stores routing tables and related information, and the other segment primarily stores firewall-filter-related information. As the allocation of more memory for routing tables or firewall filters might disrupt the forwarding operations of a Packet Forwarding Engine, the Junos OS CLI displays a warning to restart all affected FPCs when you commit a configuration that includes the memory-enhanced route statement.

The following sections describe the computation equations and the interpretation of the different memory regions for I-chip-based and Trio-based line cards:

Resource Monitoring and Usage Computation For Trio-Based Line Cards

In Trio-based line cards, memory blocks for next-hop and firewall filters are allocated separately. Also, an expansion memory is present, which is used when the allocated memory for next-hop or firewall filter is fully consumed. Both next-hop and firewall filters can allocate memory from the expansion memory. The encapsulation memory region is specific to I-chip-based line cards and it is not applicable to Trio-based line cards. Therefore, for Trio-based line cards, the percentage of free memory space can be interpreted as follows:

$$\% \text{ Free (NH)} = (1 - (\text{Used NH memory} + \text{Used Expansion memory}) / (\text{Total NH memory} + \text{Total Expansion memory})) \times 100$$
$$\% \text{ Free (Firewall or Filter)} = (1 - (\text{Used FW memory} + \text{Used Expansion memory}) / (\text{Total FW memory} + \text{Total Expansion memory})) \times 100$$

Encapsulation memory is I-chip-specific and is not applicable for Trio-based line cards.

$\% \text{ Free (Encap memory)} = \text{Not applicable}$

Resource Monitoring and Usage Computation For I-Chip-Based Line Cards

I-chip-based line cards contain 32 MB of static RAM (SRAM) memory associated with the route lookup block and 16 MB of SRAM memory associated with the output WAN block.

The route-lookup memory is a single pool of 32 MB memory that is divided into two segments of 16 MB each. In a standard configuration, segment 0 is used for NH and prefixes, and segment 1 is used for firewall or filter. This allocation can be modified by using the route-memory-enhanced option at the [edit chassis] hierarchy level. In a general configuration, NH application can be allocated memory from any of the two segments. Therefore, the percentage of free memory for NH is calculated on 32 MB memory. Currently, firewall applications are allotted memory only from segment 1. As a result, the percentage of free memory to be monitored for firewall starts from the available 16 MB memory in segment 1 only.

For I-chip-based line cards, the percentage of free memory space can be interpreted as follows:

$$\% \text{ Free (NH)} = (32 - (\text{Used NH memory} + \text{Used FW memory} + \text{Used Other application})) / 32 \times 100$$

$$\% \text{ Free (Firewall or Filter)} = (16 - (\text{Used NH memory} + \text{Used FW memory} + \text{Used Other application})) / 16 \times 100$$

The memory size for Output WAN (lwo) SRAM is 16 MB and stores the Layer 2 descriptors that contain the encapsulation information. This entity is a critical resource and needs to be monitored. This memory space is displayed in the output of the show command as "Encap mem". The percentage of free memory for the encapsulation region is calculated as follows:

$$\% \text{ Free (Encapsulation memory)} = (16 - (\text{lwo memory used (L2 descriptors + other applications)})) / 16 \times 100$$

The watermark level configured for next-hop memory is also effective for encapsulation memory. Therefore, if the percentage of free memory for encapsulation region falls below the configured watermark, logs are generated.

If the free memory percentage is lower than the free memory watermark of a specific memory type, the following error message is recorded in the syslog:

"Resource Monitor: FPC <slot no> PFE <pfe inst> <"JNH memory" or "FW/ Filter memory"> is below set watermark <configured watermark>".

You can configure resource-monitoring tracing operations by using the **traceoptions file <filename> flag flag level level size bytes** statement at the **[edit system services resource-monitor]** hierarchy level. By default, messages are written to **/var/log/rsmonlog**. The error logs associated with socket communication failure (between the Routing Engine and the Packet Forwarding Engine) are useful in diagnosing the problems in the communication between the Routing Engine and the Packet Forwarding Engine.

From the Ukern perspective, MPC5E contains only one Packet Forwarding Engine instance. The show chassis fabric plane command output displays the state of fabric plane connections to the Packet Forwarding Engine. Because two Packet Forwarding Engines exist, you notice PFE-0 and PFE-1 in the output.

```
user@host# run show chassis fabric plane
Fabric management PLANE state
Plane 0
  Plane state: ACTIVE
    FPC 0
      PFE 0 :Links ok
      PFE 1 :Links ok
```

Because only one Packet Forwarding Engine instance for MPC5E exists, the output of the `show system resource-monitor fpc` command displays only one row corresponding to Packet Forwarding Engine instance 0.

```
user@host# run show system resource-monitor fpc
FPC Resource Usage Summary
```

```
Free Heap Mem Watermark      : 20 %
Free NH Mem Watermark        : 20 %
Free Filter Mem Watermark    : 20 %
```

```
* - Watermark reached
```

	Heap		ENCAP mem	NH mem	FW
mem					
Slot #	% Free	PFE #	% Free	% Free	% Free
0	94	0		NA	83
99					

The configured watermark is retained across GRES and unified ISSU procedures.

Related •
Documentation

Diagnosing and Debugging System Performance by Configuring Memory Resource Usage Monitoring on MX Series Routers

Supported Platforms [MX104, MX2010, MX2020, MX240, MX480, MX80, MX960](#)

Junos OS supports a resource monitoring capability using both the CLI and SNMP MIB queries. You can employ this utility to provision sufficient headroom (memory space limits that are set for the application or virtual router) for ensuring system stability, especially the health and operating efficiency of I-chip-based line cards and Trio-based FPCs on MX Series routers. When the memory utilization, either the ukernel memory or ASIC memory reaches a certain threshold, the system operations compromise on the health and traffic-handling stability of the line card and such a trade-off on the system performance can be detrimental for supporting live traffic and protocols.

To configure the properties of the memory resource-utilization functionality:

1. Specify that you want to configure the monitoring mechanism for utilization of different memory resource regions.

```
[edit]
```

```
user@host# edit system services resource-monitor
```

This feature is enabled by default and you cannot disable it manually.

2. Specify the high threshold value, exceeding which warnings or error logs are generated, for all the regions of memory, such as heap or ukernel, next-hop and encapsulation, and firewall filter memory.

```
[edit system services resource-monitor]
user@host# set high-threshold value
```

3. Specify the percentage of free memory space used for next-hops to be monitored with a watermark value.

```
[edit system services resource-monitor]
user@host# set free-nh-memory-watermark percentage
```

4. Specify the percentage of free memory space used for ukernel or heap memory to be monitored with a watermark value.

```
[edit system services resource-monitor]
user@host# set free-heap-memory- watermark percentage
```

5. Specify the percentage of free memory space used for firewall and filter memory to be monitored with a watermark value.

```
[edit system services resource-monitor]
user@host# set free-filter-memory-memory- watermark percentage
```



NOTE:

The default value and the configured value of the watermark value for the percentage of free next-hop memory also applies to encapsulation memory. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20 percent.

6. Disable the generation of error log messages when the utilization of memory resources exceeds the threshold or checkpoint levels. By default, messages are written to `/var/log/rsmonlog`.

```
[edit system services resource-monitor]
user@host# set no-logging
```

7. Define the resource category that you want to monitor and analyze for ensuring system stability, especially the health and operating efficiency of I-chip-based line cards and Trio-based FPCs on MX Series routers. The resource category includes detailed CPU utilization, session rate, and session count statistics. You use the resource category statistics to understand the extent to which new attack objects or applications affect performance.

```
[edit system services resource-monitor]
user@host# edit resource-category jtree
```



NOTE: The jtree memory on all MX Series router Packet Forwarding Engines has two segments: one segment primarily stores routing tables and related information, and the other segment primarily stores firewall-filter-related information. The Junos OS provides the memory-enhanced statement to reallocate the jtree memory for routes, firewall filters, and Layer 3 VPNs.

8. Configure the type of resource as contiguous pages for which you want to enable the monitoring mechanism to provide sufficient headroom for ensuring effective system performance and traffic-handling capacity. Specify the high and low threshold value, exceeding which warnings or error logs are generated, for the specified type or region of memory, which is contiguous page in this case.

```
[edit system services resource-monitor resource-category jtree]
user@host# set resource-type contiguous-pages high-threshold percentage
user@host# set resource-type contiguous-pages low-threshold percentage
```

9. Configure the type of resource as free double words (dwords) for which you want to enable the monitoring mechanism to provide sufficient headroom for ensuring effective system performance and traffic-handling capacity. Specify the high and low threshold value, exceeding which warnings or error logs are generated, for the specified type or region of memory, which is free dwords in this case.

```
[edit system services resource-monitor resource-category jtree]
user@host# set resource-type free-dwords high-threshold percentage
user@host# set resource-type free-dwords low-threshold percentage
```

10. Configure the type of resource as free memory pages for which you want to enable the monitoring mechanism to provide sufficient headroom for ensuring effective system performance and traffic-handling capacity. Specify the high and low threshold value, exceeding which warnings or error logs are generated, for the specified type or region of memory, which is free memory pages in this case.

```
[edit system services resource-monitor resource-category jtree]
user@host# set resource-type free-pages high-threshold percentage
user@host# set resource-type free-pages low-threshold percentage
```

11. View the utilization of memory resources on the Packet Forwarding Engines of an FPC by using the **show system resource-monitor fpc** command. The filter memory denotes the filter counter memory used for firewall filter counters. The asterisk (*) displayed next to each of the memory regions denotes the ones for which the configured threshold is being currently exceeded.

```
user@host# run show system resource-monitor fpc
FPC Resource Usage Summary
```

```
Free Heap Mem Watermark      : 20 %
Free NH Mem Watermark        : 20 %
Free Filter Mem Watermark    : 20 %
```

* - Watermark reached

FW mem		Heap		ENCAP mem		NH mem	
Slot #	Free	% Free	PFE #	% Free	% Free	%	
0	99	94	0		NA	83	

Related •
Documentation

Troubleshooting the Mismatch of jnxNatObjects Values for MS-DPC and MS-MIC

Supported Platforms [MX Series](#)

Problem **Description:** When both MS-DPC and MS-MIC are deployed in a network and the Network Address Translation (NAT) type is configured as **napt-44**, the output of the **snmp mib walk** command for jnxNatObjects displays different values for MS-DPC and MS-MIC.

Resolution *Configure SNMP to Match jnxNatObjects Values for MS-DPC and MS-MIC*

To configure SNMP to match jnxNatObjects values for MS-DPC and MS-MIC:

1. Run the **set services service-set service-set-name nat-options snmp-value-match-msmic** configuration mode command. The following configuration example shows how to configure SNMP to match the values for MS-MIC-specific objects in the jnxNatObjects MIB table with the values for MS-DPC objects.

```
[edit]
user@host# set services service-set Mobile nat-options snmp-value-match-msmic
```

2. Issue the **commit** command to confirm the changes.

```
[edit]
user@host# commit
commit complete
```

3. (Optional) Run the **show snmp mib walk jnxNatObjects** command to verify that the values for MS-MIC-specific objects in the jnxNatObjects MIB table match the values for MS-DPC objects. For example, the following output shows that the values for MS-MIC-specific objects and MS-DPC objects match.

```
[edit]
user@host# run show snmp mib walk jnxNatObjects
jnxNatSrcXlatedAddrType.6.77.111.98.105.108.101 = 1
jnxNatSrcPoolType.6.77.111.98.105.108.101 = 13
jnxNatSrcNumPortAvail.6.77.111.98.105.108.101 = 64512
jnxNatSrcNumPortInuse.6.77.111.98.105.108.101 = 0
jnxNatSrcNumAddressAvail.6.77.111.98.105.108.101 = 1
jnxNatSrcNumAddressInUse.6.77.111.98.105.108.101 = 0
jnxNatSrcNumSessions.6.77.111.98.105.108.101 = 0
```

```
jnxNatRuleType.9.77.111.98.105.108.101.58.116.49 = 13
jnxNatRuleTransHits.9.77.111.98.105.108.101.58.116.49 = 0
jnxNatPoolType.6.77.111.98.105.108.101 = 13
jnxNatPoolTransHits.6.77.111.98.105.108.101 = 0
```



NOTE: You can use the delete services service-set *service-set-name* nat-options snmp-value-match-msmic configuration mode command to disable this feature.

- Related Documentation**
- *Configuring Service Rules*
 - *snmp-value-match-msmic on page 754*

Managed Objects for Ukernl Memory for a Packet Forwarding Engine in an FPC Slot

Supported Platforms MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

The **jnxPfeMemoryUkernTable**, whose object identifier is **{jnxPfeMemory 1}**, contains the **JnxPfeMemoryUkernEntry** that retrieves the global ukernel or heap memory statistics for the specified Packet Forwarding Engine slot. Each **JnxPfeMemoryUkernEntry**, whose object identifier is **{jnxPfeMemoryUkernTable 1}**, contains the objects listed in the following table. The **jnxPfeMemoryUkernEntry** denotes the memory utilization, such as the total available memory and the percentage of memory used.

Table 106: jnxPfeMemoryUkernTable

Object	Object ID	Description
jnxPfeMemoryUkernFreePercent	jnxPfeMemoryUkernEntry 3	Denotes the percentage of free Packet Forwarding Engine memory within the ukernel heap.

Managed Objects for Packet Forwarding Engine Memory Statistics Data

Supported Platforms MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

The **jnxPfeMemory** table, whose object identifier is **{jnxPfeMib 2}** contains the objects listed in [Table 107 on page 524](#)

Table 107: jnxPfeMemory Table

Object	Object ID	Description
jnxPfeMemoryUkernTable	jnxPfeMemory 1	Provides global ukernel memory statistics for the specified Packet Forwarding Engine slot.

Table 107: jnxPfeMemory Table (*continued*)

Object	Object ID	Description
jnxPfeMemoryForwardingTable	jnxPfeMemory 2	Provides global next-hop (for Trio-based line cards) or Jtree (for I-chip-based line cards) memory utilization and firewall filter memory utilization statistics for the specified Packet Forwarding Engine slot.

Managed Objects for Next-Hop, Jtree, and Firewall Filter Memory for a Packet Forwarding Engine in an FPC Slot

Supported Platforms MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

The **jnxPfeMemoryForwardingTable**, whose object identifier is **{jnxPfeMemory 2}**, contains **JnxPfeMemoryForwardingEntry** that retrieves the next-hop memory for Trio-based line cards, jtree memory for I-chip-based line cards, and firewall or filter memory statistics for the specified Packet Forwarding Engine slot for both I-chip and Trio-based line cards. Each **jnxPfeMemoryForwardingEntry**, whose object identifier is **{jnxPfeMemoryForwardingTable 1}**, contains the objects listed in the following table.

The **jnxPfeMemoryForwardingEntry** represents the ASIC instance, ASIC memory used, and ASIC free memory. The jtree memory on all MX Series router Packet Forwarding Engines has two segments: one segment primarily stores routing tables and related information, and the other segment primarily stores firewall-filter-related information. As the allocation of more memory for routing tables or firewall filters might disrupt the forwarding operations of a Packet Forwarding Engine, the Junos OS CLI displays a warning to restart all affected FPCs when you commit a configuration that includes the memory-enhanced route statement. The configuration does not become effective until you restart the FPC or DPC (on MX Series routers).

Table 108: jnxPfeMemoryForwardingTable

Object	Object ID	Description
jnxPfeMemoryForwardingChipSlot	jnxPfeMemoryForwardingEntry 1	Indicates the ASIC instance number in the Packet Forwarding Engine complex.
jnxPfeMemoryType	jnxPfeMemoryForwardingEntry 2	Indicates the Packet Forwarding Engine memory type, where nh = 1, fw = 2, encap = 3.
jnxPfeMemoryForwardingPercentFree	jnxPfeMemoryForwardingEntry 3	Indicates the percentage of memory free for each memory type.

jnxPfeMemoryErrorsTable

Supported Platforms MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

The Juniper Networks enterprise-specific Packet Forwarding Engine MIB, whose object ID is **{jnxPfeMibRoot 1}**, supports a new MIB table, **jnxPfeMemoryErrorsTable**, to display Packet Forwarding Engine memory error counters. The **jnxPfeMemoryErrorsTable**, whose

object identifier is **jnxPfeNotification 3**, contains the **JnxPfeMemoryErrorsEntry**. Each **JnxPfeMemoryErrorsEntry**, whose object identifier is **{ jnxPfeMemoryErrorsTable 1 }**, contains the objects listed in the following table.

Table 109: jnxPfeMemoryErrorsTable

Object	Object ID	Description
jnxPfeFpcSlot	jnxPfeMemoryErrorsEntry 1	Signifies the FPC slot number for this set of PFE notification
jnxPfeSlot	jnxPfeMemoryErrorsEntry 2	Signifies the PFE slot number for this set of errors
jnxPfeParityErrors	jnxPfeMemoryErrorsEntry 3	Signifies the parity error count
jnxPfeEccErrors	jnxPfeMemoryErrorsEntry 4	Signifies the error-checking code (ECC) error count

pfeMemoryErrors

Supported Platforms [MX104](#), [MX2010](#), [MX2020](#), [MX240](#), [MX480](#), [MX80](#), [MX960](#)

The **pfeMemoryErrorsNotificationPrefix**, whose object identifier is **{jnxPfeNotification 0}**, contains the **pfeMemoryErrors** attribute. The **pfeMemoryErrors** object, whose identifier is **{pfeMemoryErrorsNotificationPrefix 1}** contains the **jnxPfeParityErrors** and **jnxPfeEccErrors** objects.

Table 110: pfeMemoryErrors

Object	Object ID	Description
pfeMemoryErrors	pfeMemoryErrorsNotificationPrefix 1	A pfeMemoryErrors notification is sent when the value of jnxPfeParityErrors or jnxPfeEccErrors increases.

PART 9

Troubleshooting

- [Configuring Data Path Debugging and Trace Options on page 529](#)
- [Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits on page 547](#)
- [Using Packet Capture to Analyze Network Traffic on page 565](#)
- [Troubleshooting Security Devices on page 591](#)

CHAPTER 31

Configuring Data Path Debugging and Trace Options

- [Understanding Data Path Debugging for SRX Series Devices on page 529](#)
- [Debugging the Data Path \(CLI Procedure\) on page 530](#)
- [Example: Configuring End-to-End Debugging on SRX Series Device on page 531](#)
- [Understanding Security Debugging Using Trace Options on page 536](#)
- [Setting Security Trace Options \(CLI Procedure\) on page 537](#)
- [Displaying Log and Trace Files on page 538](#)
- [Displaying Output for Security Trace Options on page 538](#)
- [Displaying Multicast Trace Operations on page 539](#)
- [Using the J-Web Traceroute Tool on page 540](#)
- [J-Web Traceroute Results and Output Summary on page 542](#)
- [Understanding Flow Debugging Using Trace Options on page 543](#)
- [Setting Flow Debugging Trace Options \(CLI Procedure\) on page 543](#)
- [Displaying a List of Devices on page 544](#)

Understanding Data Path Debugging for SRX Series Devices

Supported Platforms [SRX Series, vSRX](#)

Data path debugging, or end-to-end debugging, support provides tracing and debugging at multiple processing units along the packet-processing path. The packet filter can be executed with minimal impact to the production system.

On an SRX Series device, a packet goes through series of events involving different components from ingress to egress processing.

With the data path debugging feature, you can trace and debug (capture packets) at different data points along the processing path. The events available in the packet-processing path are: NP ingress, load-balancing thread (LBT), jexec, packet-ordering thread (POT), and NP egress. You can also enable flow module trace if the security flow trace flag for a certain module is set.

At each event, you can specify any of the four actions (count, packet dump, packet summary, and trace). Data path debugging provides filters to define what packets to capture, and only the matched packets are traced. The packet filter can filter out packets based on logical interface, protocol, source IP address prefix, source port, destination IP address prefix, and destination port.



NOTE: Datapath debugging is not supported on SRX1500 Devices.

To enable end-to-end debugging, you must perform the following steps:

1. Define the capture file and specify the maximum capture size.
2. Define the packet filter to trace only a certain type of traffic based on the requirement.
3. Define the action profile specifying the location on the processing path from where to capture the packets (for example, LBT or NP ingress).
4. Enable the data path debugging.
5. Capture traffic.
6. Disable data path debugging.
7. View or analyze the report.



NOTE:

The packet-filtering behavior for the port and interface options is as follows:

- The packet filter traces both IPv4 and IPv6 traffic if only **port** is specified.
- The packet filter traces IPv4, IPV6, and non-IP traffic if only **interface** is specified.

**Related
Documentation**

- [Understanding Security Debugging Using Trace Options on page 536](#)
- [Understanding Flow Debugging Using Trace Options on page 543](#)
- [Debugging the Data Path \(CLI Procedure\) on page 530](#)
- [Example: Configuring End-to-End Debugging on SRX Series Device on page 531](#)

Debugging the Data Path (CLI Procedure)

Supported Platforms [SRX Series, vSRX](#)



NOTE: Datapath debugging is not supported on SRX1500 Devices.

To configure the device for data path debugging:

1. Specify the following request command to set the data path debugging for the multiple processing units along the packet-processing path:

```
[edit]
user@host# set security datapath-debug
```

2. Specify the trace options for data path-debug using the following command:

```
[edit]
user@host# set security datapath-debug traceoptions
```

3. Using the request security packet-filter command, you can set the packet filter to specify the related packets to perform data path-debug action. A maximum of four filters are supported at the same time. For example, the following command sets the first packet-filter:

```
[edit]
user@host# set security datapath-debug packet-filter name
```

4. Using the request security action-profile command, you can set the action for the packet match for a specified filter. Only the default action profile is supported, which is the trace option for network processor ezchip ingress, ezchip egress, spu.lbt, and spu.pot:

```
[edit]
user@host# set security datapath-debug packet-filter name action-profile
```

Related Documentation

- [Understanding Data Path Debugging for SRX Series Devices on page 529](#)
- [Understanding Security Debugging Using Trace Options on page 536](#)
- [Understanding Flow Debugging Using Trace Options on page 543](#)
- [Setting Flow Debugging Trace Options \(CLI Procedure\) on page 543](#)

Example: Configuring End-to-End Debugging on SRX Series Device

Supported Platforms SRX5400, SRX5600, SRX5800

This example shows how to configure and enable end-to-end debugging on an SRX Series device with an SRX5K-MPC.

- [Requirements on page 532](#)
- [Overview on page 532](#)
- [Configuration on page 532](#)

- [Enabling Data Path Debugging on page 535](#)
- [Verification on page 535](#)

Requirements

This example uses the following hardware and software components:

- SRX5600 device with an SRX5K-MPC installed with 100-Gigabit Ethernet CFP transceiver
- Junos OS Release 12.1X47-D15 or later for SRX Series devices

Before you begin:

- See *Understanding Data Path Debugging for SRX Series Devices*.

No special configuration beyond device initialization is required before configuring this feature.

Overview

Data path debugging enhances troubleshooting capabilities by providing tracing and debugging at multiple processing units along the packet-processing path. With the data path debugging feature, you can trace and debug (capture packets) at different data points along the processing path. At each event, you can specify an action (count, packet dump, packet summary, and trace) and you can set filters to define what packets to capture.

In this example, you define a traffic filter, and then you apply an action profile. The action profile specifies a variety of actions on the processing unit. The ingress and egress are specified as locations on the processing path to capture the data for incoming and outgoing traffic.

Next, you enable data path debugging in operational mode, and finally you view the data capture report.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security datapath-debug traceoptions file e2e.trace size 10m
set security datapath-debug capture-file datapcap format pcap
set security datapath-debug maximum-capture-size 1500
set security datapath-debug action-profile profile-1 preserve-trace-order
set security datapath-debug action-profile profile-1 record-pic-history
set security datapath-debug action-profile profile-1 event np-ingress trace
set security datapath-debug action-profile profile-1 event np-ingress count
set security datapath-debug action-profile profile-1 event np-ingress packet-summary
set security datapath-debug action-profile profile-1 event np-ingress packet-count
set security datapath-debug action-profile profile-1 event np-egress trace
set security datapath-debug action-profile profile-1 event np-egress count
```

```

set security datapath-debug action-profile profile-1 event np-egress packet-summary
set security datapath-debug action-profile profile-1 event np-egress packet-count
set security datapath-debug packet-filter filter-1
set security datapath-debug packet-filter filter-1 action-profile profile-1
set security datapath-debug packet-filter filter-1 protocol tcp
set security datapath-debug packet-filter filter-1 source-prefix 203.0.113.1/24
set security datapath-debug packet-filter filter-1 destination-prefix 203.0.113.4/24
set security datapath-debug packet-filter filter-1 source-port 1000
set security datapath-debug packet-filter filter-1 destination-port 80
set security datapath-debug packet-filter filter-1 interface xe-2/2/0.0

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure data path debugging:

1. Edit the security datapath debugging option for the multiple processing units along the packet-processing path:

```

[edit]
user@host# edit security datapath-debug

```

2. Enable the capture file, file format, file size, and the number of files.

```

[edit security datapath-debug]
user@host# set traceoptions file e2e.trace size 10m
user@host# set capture-file datapcap format pcap;
user@host# set maximum-capture-size 1500

```

3. Configure action profile, event type, and actions for the action profile.

```

[edit security datapath-debug]
user@host# set action-profile profile-1 preserve-trace-order
user@host# set action-profile profile-1 record-pic-history
user@host# set action-profile profile-1 event np-ingress trace
user@host# set action-profile profile-1 event np-ingress count
user@host# set action-profile profile-1 event np-ingress packet-summary
user@host# set action-profile profile-1 event np-ingress packet-count
user@host# set action-profile profile-1 event np-egress trace
user@host# set action-profile profile-1 event np-egress count
user@host# set action-profile profile-1 event np-egress packet-summary
user@host# set action-profile profile-1 event np-egress packet-count

```

4. Configure packet filter, action, and filter options.

```

[edit security datapath-debug]
user@host# set packet-filter filter-1
user@host# set packet-filter filter-1 action-profile profile-1
user@host# set packet-filter filter-1 protocol tcp
user@host# set packet-filter filter-1 source-prefix 203.0.113.1/24
user@host# set packet-filter filter-1 destination-prefix 203.0.113.4/24
user@host# set packet-filter filter-1 source-port 1000
user@host# set packet-filter filter-1 destination-port 80

```

```
user@host# set packet-filter filter-1 interface xe-2/2/0.0
```



NOTE: You must configure multiple packet-filter statements to capture the traffic, because one packet-filter captures only the traffic that is specified in it, and the same packet filter does not capture the traffic in the reverse direction. You must specify the source and destination address for each packet filter, else the SRX Series device might log the entire traffic and cause adverse impact to traffic.

Results From configuration mode, confirm your configuration by entering the **show security datapath-debug** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
traceoptions {
  file e2e.trace size 10m;
}
capture-file datapcap format pcap;
maximum-capture-size 1500;
action-profile {
  profile-1 {
    preserve-trace-order;
    record-pic-history;
    event np-ingress {
      trace;
      count;
      packet-summary;
      packet-dump;
    }
    event np-egress {
      trace;
      count;
      packet-summary;
      packet-dump;
    }
  }
}
packet-filter filter-1 {
  action-profile profile-1;
  protocol tcp;
  source-prefix 203.0.113.1/24;
  destination-prefix 203.0.113.4/24;
  source-port 1000;
  destination-port 80;
  interface xe-2/2/0.0;
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Enabling Data Path Debugging

- Step-by-Step Procedure** After configuring data path debugging, you must start the process on the device from operational mode.
1. Enable data path debugging.

```
user@host> request security datapath-debug capture start
```

datapath-debug capture started on file datapcap
 2. Before you verify the configuration and view the reports, you must disable data path debugging.

```
user@host> request security datapath-debug capture stop
```

datapath-debug capture successfully stopped, use show security datapath-debug capture to view



NOTE: You must stop the debug process after you have finished capturing the data. If you attempt to open the captured files without stopping the debug process, the files obtained cannot be opened through any third-party software (for example, tcpdump and wireshark).

Verification

Confirm that the configuration is working properly.

Verifying Data Path Debug Packet Capture Details

Purpose Verify the data captured by enabling the data path debugging configuration.

Action From operational mode, enter the **show security datapath-debug capture** command.

```
Packet 8, len 152: (C2/F2/P0/SEQ:57935:np-ingress)
00 10 db ff 10 02 00 30 48 83 8d 4f 08 00 45 00
00 54 00 00 40 00 40 01 9f c7 c8 07 05 69 c8 08
05 69 08 00 91 1f 8f 03 2a a2 ae 66 85 53 8c 7d
02 00 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15
16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25
26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35
36 37
Packet 9, len 152: (C2/F2/P0/SEQ:57935:np-egress)
00 30 48 8d 1a bf 00 10 db ff 10 03 08 00 45 00
00 54 00 00 40 00 3f 01 a0 c7 c8 07 05 69 c8 08
05 69 08 00 91 1f 8f 03 2a a2 ae 66 85 53 8c 7d
02 00 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15
16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25
26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35
36 37....
```

For brevity, the **show** command output is truncated to display only a few samples. Additional samples have been replaced with ellipses (...).

To view the results, from CLI operational mode, access the local UNIX shell and navigate to the directory `/var/log/<file-name>`. The result can be read by using the **tcpdump** utility.

```
user@host>start shell
%tcpdump -nr/var/log/e2e.pcap
```

```
21:50:04.288767 C0/F3 event:1(np-ingress) SEQ:1 IP 192.168.14.2 > 192.168.13.2:
ICMP echo request, id 57627, seq 0, length 64
21:50:04.292590 C0/F3 event:2(np-egress) SEQ:1 IP 192.168.14.2 > 192.168.13.2:
ICMP echo request, id 57627, seq 0, length 64
1:50:04.295164 C0/F3 event:1(np-ingress) SEQ:2 IP 192.168.13.2 > 192.168.14.2:
ICMP echo reply, id 57627, seq 0, length 64
21:50:04.295284 C0/F3 event:2(np-egress) SEQ:2 IP 192.168.13.2 > 192.168.14.2:
ICMP echo reply, id 57627, seq 0, length 64
```



NOTE: If you are finished with troubleshooting the data path debugging, remove all traceoptions (not limited to flow traceoptions) and the complete data path debug configuration, including the data path debug configuration for packet capturing (packet-dump), which needs to be started/stopped manually. If any part of the debugging configuration remains active, it will continue to use the resources of the device (CPU/memory).

Related Documentation • [Understanding Data Path Debugging for SRX Series Devices on page 529](#)

Understanding Security Debugging Using Trace Options

Supported Platforms [SRX Series, vSRX](#)

The Junos OS trace function allows applications to write security debugging information to a file. The information that appears in this file is based on criteria you set. You can use this information to analyze security application issues.

The trace function operates in a distributed manner, with each thread writing to its own trace buffer. These trace buffers are then collected at one point, sorted, and written to trace files. Trace messages are delivered using the InterProcess Communications (IPC) protocol. A trace message has a lower priority than that of control protocol packets such as BGP, OSPF, and IKE, and therefore delivery is not considered to be as reliable.

Related Documentation • [Understanding Data Path Debugging for SRX Series Devices on page 529](#)
• [Understanding Flow Debugging Using Trace Options on page 543](#)
• [Setting Security Trace Options \(CLI Procedure\) on page 537](#)
• [Debugging the Data Path \(CLI Procedure\) on page 530](#)

- [Displaying Output for Security Trace Options on page 538](#)

Setting Security Trace Options (CLI Procedure)

Supported Platforms [SRX Series, vSRX](#)

Use the following configuration statements to configure security trace options in the CLI configuration editor.

- To disable remote tracing, enter the following statement:

```
[edit]
user@host# set security traceoptions no-remote-trace
```

- To write trace messages to a local file, enter the following statement. The system saves the trace file in the `/var/log/` directory.

```
[edit]
user@host# set security traceoptions use-local-files
```

- To specify a name for the trace file, enter the following statement. Valid values range from 1 and 1024 characters. The name cannot include spaces, `/`, or `%` characters. The default filename is `security`.

```
[edit]
user@host# set security traceoptions file filename
```

- To specify the maximum number of trace files that can accumulate, enter the following statement. Valid values range from 2 to 1000. The default value is 3.

```
[edit]
user@host# set security traceoptions file files 3
```

- To specify the match criteria that you want the system to use when logging information to the file, enter the following statement. Enter a regular expression. Wildcard `(*)` characters are accepted.

```
[edit]
user@host# set security traceoptions file match *thread
```

- To allow any user to read the trace file, enter the **world-readable** statement. Otherwise, enter the **no-world-readable** statement.

```
[edit]
user@host# set security traceoptions file world-readable
user@host# set security traceoptions file no-world-readable
```

- To specify the maximum size to which the trace file can grow, enter the following statement. Once the file reaches the specified size, it is compressed and renamed `filename0.gz`, the next file is named `filename1.gz`, and so on. Valid values range from 10240 to 1,073,741,824.

```
[edit]
user@host# set security traceoptions file size 10240
```

- To turn on trace options and to perform more than one tracing operation, set the following flags.

```
[edit]
```

```
user@host# set security traceoptions flag all
user@host# set security traceoptions flag compilation
user@host# set security traceoptions flag configuration
user@host# set security traceoptions flag routing-socket
```

- To specify the groups that these trace option settings do or do not apply to, enter the following statements:

```
[edit]
user@host# set security traceoptions apply-groups value
user@host# set security traceoptions apply-groups-except value
```

**Related
Documentation**

- [Understanding Security Debugging Using Trace Options on page 536](#)
- [Displaying Output for Security Trace Options on page 538](#)

Displaying Log and Trace Files

Supported Platforms [SRX Series, vSRX](#)

Enter the **monitor start** command to display real-time additions to system logs and trace files:

```
user@host> monitor start filename
```

When the device adds a record to the file specified by *filename*, the record displays on the screen. For example, if you have configured a system log file named **system-log** (by including the **syslog** statement at the **[edit system]** hierarchy level), you can enter the **monitor start system-log** command to display the records added to the system log.

To display a list of files that are being monitored, enter the **monitor list** command. To stop the display of records for a specified file, enter the **monitor stop filename** command.

**Related
Documentation**

- [Displaying a List of Devices on page 544](#)
- [Displaying Real-Time Monitoring Information on page 389](#)

Displaying Output for Security Trace Options

Supported Platforms [SRX Series, vSRX](#)

Purpose Display output for security trace options.

Action Use the **show security traceoptions** command to display the output of your trace files. For example:

```
[edit]
user@host # show security traceoptions file usp_trace
user@host # show security traceoptions flag all
user@host # show security traceoptions rate-limit 888
```

The output for this example is as follows:

```
Apr 11 16:06:42 21:13:15.750395:CID-906489336:FPC-01:PIC-01:THREAD_ID-01:PFE:now
update 0x3607edf8df8in 0x3607e8d0
Apr 11 16:06:42 21:13:15.874058:CID-1529687608:FPC-01:PIC-01:THREAD_ID-01:CTRL:Enter
Function[util_ssam_handler]
Apr 11 16:06:42 21:13:15.874485:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default1: Rate
limit changed to 888
Apr 11 16:06:42 21:13:15.874538:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default1:
Destination ID set to 1
Apr 11 16:06:42 21:13:15.874651:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default2: Rate
limit changed to 888
Apr 11 16:06:42 21:13:15.874832:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default2:
Destination ID set to 1
Apr 11 16:06:42 21:13:15.874942:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default3: Rate
limit changed to 888
Apr 11 16:06:42 21:13:15.874997:CID-00:FPC-01:PIC-01:THREAD_ID-01:CTRL:default3:
Destination ID set to 1
```

Related Documentation

- [Understanding Security Debugging Using Trace Options on page 536](#)
- [Setting Security Trace Options \(CLI Procedure\) on page 537](#)

Displaying Multicast Trace Operations

Supported Platforms [SRX Series, vSRX](#)

To monitor and display multicast trace operations, enter the **mtrace monitor** command:

```
user@host> mtrace monitor
```

```
Mtrace query at Apr 21 16:00:54 by 192.1.30.2, resp to 224.0.1.32, qid 2a83aa
packet from 192.1.30.2 to 224.0.0.2 from 192.1.30.2 to 192.1.4.1 via group
224.1.1.1 (mxhop=60) Mtrace query at Apr 21 16:00:57 by 192.1.30.2, resp to
224.0.1.32, qid 25dc17 packet from 192.1.30.2 to 224.0.0.2 from 192.1.30.2 to
192.1.4.1 via group 224.1.1.1 (mxhop=60) Mtrace query at Apr 21 16:01:00 by
192.1.30.2, resp to same, qid 20e046 packet from 192.1.30.2 to 224.0.0.2 from
192.1.30.2 to 192.1.4.1 via group 224.1.1.1 (mxhop=60) Mtrace query at Apr 21
16:01:10 by 192.1.30.2, resp to same, qid 1d25ad packet from 192.1.30.2 to
224.0.0.2 from 192.1.30.2 to 192.1.4.1 via group 224.1.1.1 (mxhop=60)
```

This example displays only **mtrace** queries. However, when the device captures an **mtrace** response, the display is similar, but the complete **mtrace** response also appears (exactly as it is appears in the **mtrace from-source** command output).

[Table 111 on page 539](#) summarizes the output fields of the display.

Table 111: CLI mtrace monitor Command Output Summary

Field	Description
Mtrace operation-type at time-of-day	<ul style="list-style-type: none"> • operation-type—Type of multicast trace operation: query or response. • time-of-day—Date and time the multicast trace query or response was captured.

Table 111: CLI mtrace monitor Command Output Summary (*continued*)

Field	Description
by	IP address of the host issuing the query.
resp to <i>address</i>	<i>address</i> —Response destination address.
qid <i>qid</i>	<i>qid</i> —Query ID number.
packet from <i>source</i> to <i>destination</i>	<ul style="list-style-type: none"> <i>source</i>—IP address of the source of the query or response. <i>destination</i>—IP address of the destination of the query or response.
from <i>source</i> to <i>destination</i>	<ul style="list-style-type: none"> <i>source</i>—IP address of the multicast source. <i>destination</i>—IP address of the multicast destination.
via group <i>address</i>	<i>address</i> —Group address being traced.
mxhop=<i>number</i>	<i>number</i> —Maximum hop setting.

- Related Documentation**
- [Using the J-Web Traceroute Tool on page 540](#)
 - [J-Web Traceroute Results and Output Summary on page 542](#)

Using the J-Web Traceroute Tool

Supported Platforms [SRX Series, vSRX](#)

You can use the traceroute diagnostic tool to display a list of devices between the device and a specified destination host. The output is useful for diagnosing a point of failure in the path from the device to the destination host, and addressing network traffic latency and throughput problems.

The device generates the list of devices by sending a series of ICMP traceroute packets in which the time-to-live (TTL) value in the messages sent to each successive device is incremented by 1. (The TTL value of the first traceroute packet is set to 1.) In this manner, each device along the path to the destination host replies with a Time Exceeded packet from which the source IP address can be obtained.

To use the traceroute tool:

1. Select **Troubleshoot>Traceroute**.
2. Next to Advanced options, click the expand icon.
3. Enter information into the Traceroute page (see [Table 112 on page 541](#)).

Table 112: Traceroute Field Summary

Field	Function	Your Action
Remote Host	Identifies the destination host of the traceroute. The Remote Host field is the only required field.	Type the hostname or IP address of the destination host.
Advanced Options		
Don't Resolve Addresses	Determines whether hostnames of the hops along the path are displayed, in addition to IP addresses.	<ul style="list-style-type: none"> • Suppress the display of the hop hostnames by selecting the check box. • Display the hop hostnames by clearing the check box.
Gateway	Specifies the IP address of the gateway to route through.	Type the gateway IP address.
Source Address	Specifies the source address of the outgoing traceroute packets.	Type the source IP address.
Bypass Routing	<p>Determines whether traceroute packets are routed by means of the routing table.</p> <p>If the routing table is not used, traceroute packets are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, traceroute responses are not sent.</p>	<ul style="list-style-type: none"> • Bypass the routing table and send the traceroute packets to hosts on the specified interface only by selecting the check box. • Route the traceroute packets by means of the routing table by clearing the check box.
Interface	Specifies the interface on which the traceroute packets are sent.	Select the interface on which traceroute packets are sent from the list. If you select any , the traceroute requests are sent on all interfaces.
Time-to-Live	Specifies the maximum time-to-live (TTL) hop count for the traceroute request packet.	Select the TTL from the list.
Type-of-Service	Specifies the type-of-service (TOS) value to include in the IP header of the traceroute request packet.	Select the decimal value of the TOS field from the list.
Resolve AS Numbers	Determines whether the autonomous system (AS) number of each intermediate hop between the device and the destination host is displayed.	<ul style="list-style-type: none"> • Display the AS numbers by selecting the check box. • Suppress the display of the AS numbers by clearing the check box.

4. Click **Start**.

The results of the traceroute operation are displayed in the main pane. If no options are specified, each line of the traceroute display is in the following format:

hop-number host (ip-address) [as-number]time1 time2 time3

The device sends a total of three traceroute packets to each router along the path and the round-trip time for each traceroute operation appears. If the device times out

before receiving a **Time Exceeded** message, an asterisk (*) appears for that round-trip time.

5. You can stop the traceroute operation before it is complete by clicking **OK** while the results of the traceroute operation appear.

Related Documentation

- [Diagnostic Tools Overview on page 8](#)
- [J-Web Traceroute Results and Output Summary on page 542](#)
- [Using the J-Web Ping MPLS Tool on page 555](#)
- [Using the J-Web Ping Host Tool on page 552](#)
- [Using the J-Web Packet Capture Tool on page 586](#)

J-Web Traceroute Results and Output Summary

Supported Platforms [SRX Series, vSRX](#)

[Table 113 on page 542](#) summarizes the output in the traceroute display.

Table 113: J-Web Traceroute Results and Output Summary

Field	Description
<i>hop-number</i>	Number of the hop (device) along the path.
<i>host</i>	Hostname, if available, or IP address of the device. If the Don't Resolve Addresses check box is selected, the hostname does not appear.
<i>ip-address</i>	IP address of the device.
<i>as-number</i>	AS number of the device.
<i>time1</i>	Round-trip time between the sending of the first traceroute packet and the receiving of the corresponding Time Exceeded packet from that particular device.
<i>time2</i>	Round-trip time between the sending of the second traceroute packet and the receiving of the corresponding Time Exceeded packet from that particular device.
<i>time3</i>	Round-trip time between the sending of the third traceroute packet and the receiving of the corresponding Time Exceeded packet from that particular device.

If the device does not display the complete path to the destination host, one of the following explanations might apply:

- The host is not operational.
- There are network connectivity problems between the device and the host.
- The host, or a router along the path, might be configured to ignore ICMP traceroute messages.

- The host, or a device along the path, might be configured with a firewall filter that blocks ICMP traceroute requests or ICMP time exceeded responses.
- The value you selected in the Time Exceeded box was less than the number of hops in the path to the host. In this case, the host might reply with an ICMP error message.

- Related Documentation**
- [Diagnostic Tools Overview on page 8](#)
 - [Using the J-Web Traceroute Tool on page 540](#)

Understanding Flow Debugging Using Trace Options

Supported Platforms [SRX Series, vSRX](#)

For flow trace options, you can define a packet filter using combinations of **destination-port**, **destination-prefix**, **interface**, **protocol**, **source-port**, and **source-prefix**. If the security flow trace flag for a certain module is set, the packet matching the specific packet filter triggers flow tracing and writes debugging information to the trace file.

- Related Documentation**
- [Understanding Data Path Debugging for SRX Series Devices on page 529](#)
 - [Understanding Security Debugging Using Trace Options on page 536](#)
 - [Setting Flow Debugging Trace Options \(CLI Procedure\) on page 543](#)
 - [Debugging the Data Path \(CLI Procedure\) on page 530](#)

Setting Flow Debugging Trace Options (CLI Procedure)

Supported Platforms [SRX Series](#)

The following examples display the options you can set by using **security flow traceoptions**.

- To match the imap destination port for the filter1 packet filter, use the following statement:

```
[edit]
user@host# set security flow traceoptions packet-filter filter1 destination-port imap
```
- To set the 1.2.3.4 destination IPv4 prefix address for the filter1 packet filter, use the following statement:

```
[edit]
user@host# set security flow traceoptions packet-filter filter1 destination-prefix 1.2.3.4
```
- To set the fxp0 logical interface for the filter1 packet filter, use the following statement:

```
[edit]
user@host# set security flow traceoptions packet-filter filter1 interface fxp0
```
- To match the TCP IP protocol for the filter1 packet filter, use the following statement:

```
[edit]
user@host# set security flow traceoptions packet-filter filter1 protocol tcp
```

- To match the HTTP source port for the filter1 packet filter, use the following statement:

```
[edit]
user@host# set security flow traceoptions packet-filter filter1 source-port http
```

- To set the 5.6.7.8 IPv4 prefix address for the filter1 packet filter, use the following statement:

```
[edit]
user@host# set security flow traceoptions packet-filter filter1 source-prefix 5.6.7.8
```

Related Documentation

- [Understanding Flow Debugging Using Trace Options on page 543](#)

Displaying a List of Devices

Supported Platforms [SRX Series, vSRX](#)

To display a list of devices between the device and a specified destination host, enter the **traceroute** command with the following syntax:

```
user@host> traceroute host <interface interface-name> <as-number-lookup>
<bypass-routing> <gateway address> <inet | inet6> <no-resolve>
<routing-instance routing-instance-name> <source source-address> <tos number>
<tll number> <wait seconds>
```

[Table 114 on page 544](#) describes the **traceroute** command options.

Table 114: CLI traceroute Command Options

Option	Description
<i>host</i>	Sends traceroute packets to the hostname or IP address you specify.
<i>interface interface-name</i>	(Optional) Sends the traceroute packets on the interface you specify. If you do not include this option, traceroute packets are sent on all interfaces.
<i>as-number-lookup</i>	(Optional) Displays the autonomous system (AS) number of each intermediate hop between the device and the destination host.
<i>bypass-routing</i>	(Optional) Bypasses the routing tables and sends the traceroute packets only to hosts on directly attached interfaces. If the host is not on a directly attached interface, an error message is returned. Use this option to display a route to a local system through an interface that has no route through it.
<i>gateway address</i>	(Optional) Uses the gateway you specify to route through.
<i>inet</i>	(Optional) Forces the traceroute packets to an IPv4 destination.
<i>inet6</i>	(Optional) Forces the traceroute packets to an IPv6 destination.
<i>no-resolve</i>	(Optional) Suppresses the display of the hostnames of the hops along the path.

Table 114: CLI traceroute Command Options (*continued*)

Option	Description
routing-instance <i>routing-instance-name</i>	(Optional) Uses the routing instance you specify for the traceroute.
source address	(Optional) Uses the source address that you specify, in the traceroute packet.
tos number	(Optional) Sets the type-of-service (TOS) value in the IP header of the traceroute packet. Specify a value from 0 through 255.
ttl number	(Optional) Sets the time-to-live (TTL) value for the traceroute packet. Specify a hop count from 0 through 128.
wait seconds	(Optional) Sets the maximum time to wait for a response.

To quit the **traceroute** command, press Ctrl-C.

The following is sample output from a **traceroute** command:

```
user@host> traceroute host2
```

```
traceroute to 173.24.232.66 (172.24.230.41), 30 hops max, 40 byte packets 1
173.18.42.253 (173.18.42.253) 0.482 ms 0.346 ms 0.318 ms 2 host4.site1.net
(173.18.253.5) 0.401 ms 0.435 ms 0.359 ms 3 host5.site1.net (173.18.253.5)
0.401 ms 0.360 ms 0.357 ms 4 173.24.232.65 (173.24.232.65) 0.420 ms 0.456
ms 0.378 ms 5 173.24.232.66 (173.24.232.66) 0.830 ms 0.779 ms 0.834 ms
```

The fields in the display are the same as those displayed by the J-Web traceroute diagnostic tool.

Related Documentation

- [Displaying Log and Trace Files on page 538](#)

CHAPTER 32

Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits

- [MPLS Connection Checking Overview on page 547](#)
- [Understanding Ping MPLS on page 549](#)
- [Using the ping Command on page 550](#)
- [Using the J-Web Ping Host Tool on page 552](#)
- [J-Web Ping Host Results and Output Summary on page 554](#)
- [Using the J-Web Ping MPLS Tool on page 555](#)
- [J-Web Ping MPLS Results and Output Summary on page 558](#)
- [Pinging Layer 2 Circuits on page 559](#)
- [Pinging Layer 2 VPNs on page 560](#)
- [Pinging Layer 3 VPNs on page 562](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 563](#)

MPLS Connection Checking Overview

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

Use either the J-Web ping MPLS diagnostic tool or the CLI commands **ping mpls**, **ping mpls l2circuit**, **ping mpls l2vpn**, and **ping mpls l3vpn** to diagnose the state of label-switched paths (LSPs), Layer 2 and Layer 3 virtual private networks (VPNs), and Layer 2 circuits.

Based on how the LSP or VPN outbound (egress) node at the remote endpoint of the connection replies to the probes, you can determine the connectivity of the LSP or VPN.

Each probe is an echo request sent to the LSP or VPN exit point as an MPLS packet with a UDP payload. If the outbound node receives the echo request, it checks the contents of the probe and returns a value in the UDP payload of the response packet. If the device receives the response packet, it reports a successful ping response.

Responses that take longer than 2 seconds are identified as failed probes.

[Table 115 on page 548](#) summarizes the options for using either the J-Web ping MPLS diagnostic tool or the CLI **ping mpls** command to display information about MPLS connections in VPNs and LSPs.

Table 115: Options for Checking MPLS Connections

J-Web Ping MPLS Tool	ping mpls Command	Purpose	Additional Information
Ping RSVP-signaled LSP	ping mpls rsvp	Checks the operability of an LSP that has been set up by the Resource Reservation Protocol (RSVP). The device pings a particular LSP using the configured LSP name.	When an RSVP-signaled LSP has several paths, the device sends the ping requests on the path that is currently active.
Ping LDP-signaled LSP	ping mpls ldp	Checks the operability of an LSP that has been set up by the Label Distribution Protocol (LDP). The device pings a particular LSP using the forwarding equivalence class (FEC) prefix and length.	When an LDP-signaled LSP has several gateways, the device sends the ping requests through the first gateway. Ping requests sent to LDP-signaled LSPs use only the master routing instance.
Ping LSP to Layer 3 VPN prefix	ping mpls l3vpn	Checks the operability of the connections related to a Layer 3 VPN. The device tests whether a prefix is present in a provider edge (PE) device's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix.	The device does not test the connection between a PE device and a customer edge (CE) router.
Locate LSP using interface name	ping mpls l2vpn interface	Checks the operability of the connections related to a Layer 2 VPN. The device directs outgoing request probes out the specified interface.	—
Instance to which this connection belongs	ping mpls l2vpn instance	Checks the operability of the connections related to a Layer 2 VPN. The device pings on a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier, to test the integrity of the Layer 2 VPN circuit (specified by the identifiers) between the inbound and outbound PE routers.	—
Locate LSP from interface name	ping mpls l2circuit interface	Checks the operability of the Layer 2 circuit connections. The device directs outgoing request probes out the specified interface.	—
Locate LSP from virtual circuit information	ping mpls l2circuit virtual-circuit	Checks the operability of the Layer 2 circuit connections. The device pings on a combination of the IPv4 prefix and the virtual circuit identifier on the outbound PE router, testing the integrity of the Layer 2 circuit between the inbound and outbound PE routers.	—

Table 115: Options for Checking MPLS Connections (*continued*)

J-Web Ping MPLS Tool	ping mpls Command	Purpose	Additional Information
Ping end point of LSP	ping mpls lsp-end-point	Checks the operability of an LSP endpoint. The device pings an LSP endpoint using either an LDP FEC prefix or an RSVP LSP endpoint address.	—

- Related Documentation**
- [Diagnostic Tools Overview on page 8](#)
 - [Understanding Ping MPLS on page 549](#)
 - [Using the J-Web Ping Host Tool on page 552](#)
 - [Using the ping Command on page 550](#)

Understanding Ping MPLS

Supported Platforms [SRX100, SRX110, SRX210, SRX220, SRX240, SRX650](#)

Before using the ping MPLS feature, make sure that the receiving interface on the VPN or LSP remote endpoint has MPLS enabled, and that the loopback interface on the outbound node is configured as **127.0.0.1**. The source address for MPLS probes must be a valid address on the J Series device.

This section includes the following topics:

- [MPLS Enabled on page 549](#)
- [Loopback Address on page 549](#)
- [Source Address for Probes on page 550](#)

MPLS Enabled

To process ping MPLS requests, the remote endpoint of the VPN or LSP must be configured appropriately. You must enable MPLS on the receiving interface of the outbound node for the VPN or LSP. If MPLS is not enabled, the remote endpoint drops the incoming request packets and returns an “ICMP host unreachable” message to the J Series device.

Loopback Address

The loopback address (**lo0**) on the outbound node must be configured as **127.0.0.1**. If this interface address is not configured correctly, the outbound node does not have this forwarding entry. It drops the incoming request packets and returns a “host unreachable” message to the J Series device.

Source Address for Probes

The source IP address you specify for a set of probes must be an address configured on one of the J Series device interfaces. If it is not a valid J Series device address, the ping request fails with the error message “Can't assign requested address.”

Related Documentation

- [Diagnostic Tools Overview on page 8](#)
- [MPLS Connection Checking Overview on page 547](#)
- [Using the J-Web Ping Host Tool on page 552](#)
- [Using the J-Web Ping MPLS Tool on page 555](#)
- [Using the ping Command on page 550](#)
- [Junos OS Interfaces Configuration Guide for Security Devices](#)
- [Junos OS Feature Support Reference for SRX Series and J Series Devices](#)

Using the ping Command

Supported Platforms [SRX Series, vSRX](#)

You can perform certain tasks only through the CLI. Use the CLI **ping** command to verify that a host can be reached over the network. This command is useful for diagnosing host and network connectivity problems. The device sends a series of ICMP echo (ping) requests to a specified host and receives ICMP echo responses.

Enter the **ping** command with the following syntax:

```
user@host> ping host <interface source-interface> <bypass-routing> <count number>
<do-not-fragment> <inet | inet6> <interval seconds> <loose-source [hosts]>
<no-resolve> <pattern string> <rapid> <record-route>
<routing-instance routing-instance-name> <size bytes> <source source-address> <strict>
<strict-source [hosts]> <tos number> <tll number> <wait seconds> <detail> <verbose>
```

[Table 116 on page 550](#) describes the **ping** command options.

To quit the **ping** command, press Ctrl-C.

Table 116: CLI ping Command Options

Option	Description
<i>host</i>	Pings the hostname or IP address you specify.
<i>interface source-interface</i>	(Optional) Sends the ping requests on the interface you specify. If you do not include this option, ping requests are sent on all interfaces.
<i>bypass-routing</i>	(Optional) Bypasses the routing tables and sends the ping requests only to hosts on directly attached interfaces. If the host is not on a directly attached interface, an error message is returned. Use this option to ping a local system through an interface that has no route through it.

Table 116: CLI ping Command Options (*continued*)

Option	Description
count <i>number</i>	(Optional) Limits the number of ping requests to send. Specify a count from 1 through 2,000,000,000. If you do not specify a count, ping requests are continuously sent until you press Ctrl-C.
do-not-fragment	(Optional) Sets the Don't Fragment (DF) bit in the IP header of the ping request packet.
inet	(Optional) Forces the ping requests to an IPv4 destination.
inet6	(Optional) Forces the ping requests to an IPv6 destination.
interval <i>seconds</i>	(Optional) Sets the interval between ping requests, in seconds. Specify an interval from 0.1 through 10,000. The default value is 1 second.
loose-source [<i>hosts</i>]	(Optional) For IPv4, sets the loose source routing option in the IP header of the ping request packet.
no-resolve	(Optional) Suppresses the display of the hostnames of the hops along the path.
pattern <i>string</i>	(Optional) Includes the hexadecimal string you specify, in the ping request packet.
rapid	(Optional) Sends ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the count option.
record-route	(Optional) For IPv4, sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed on the screen.
routing-instance <i>routing-instance-name</i>	(Optional) Uses the routing instance you specify for the ping request.
size <i>bytes</i>	(Optional) Sets the size of the ping request packet. Specify a size from 0 through 65,468. The default value is 56 bytes, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.
source <i>source-address</i>	(Optional) Uses the source address that you specify, in the ping request packet.
strict	(Optional) For IPv4, sets the strict source routing option in the IP header of the ping request packet.
strict-source [<i>hosts</i>]	(Optional) For IPv4, sets the strict source routing option in the IP header of the ping request packet, and uses the list of hosts you specify for routing the packet.
tos <i>number</i>	(Optional) Sets the type-of-service (TOS) value in the IP header of the ping request packet. Specify a value from 0 through 255.
ttl <i>number</i>	(Optional) Sets the time-to-live (TTL) value for the ping request packet. Specify a value from 0 through 255.

Table 116: CLI ping Command Options (*continued*)

Option	Description
wait <i>seconds</i>	(Optional) Sets the maximum time to wait after sending the last ping request packet. If you do not specify this option, the default delay is 10 seconds. If you use this option without the count option, the device uses a default count of 5 packets.
detail	(Optional) Displays the interface on which the ping response was received.
verbose	(Optional) Displays detailed output.

The following is sample output from a **ping** command:

```
user@host> ping host3 count 4
```

```
PING host3.site.net (176.26.232.111): 56 data bytes 64 bytes from 176.26.232.111:
icmp_seq=0 ttl=122 time=0.661 ms 64 bytes from 176.26.232.111: icmp_seq=1 ttl=122
time=0.619 ms 64 bytes from 176.26.232.111: icmp_seq=2 ttl=122 time=0.621 ms 64
bytes from 176.26.232.111: icmp_seq=3 ttl=122 time=0.634 ms --- host3.site.net
ping statistics --- 4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.619/0.634/0.661/0.017 ms
```

The fields in the display are the same as those displayed by the J-Web ping host diagnostic tool.

Related Documentation

- [Diagnostic Tools Overview on page 8](#)
- [Understanding Ping MPLS on page 549](#)
- [Pinging Layer 2 Circuits on page 559](#)
- [Pinging Layer 2 VPNs on page 560](#)
- [Pinging Layer 3 VPNs on page 562](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 563](#)
- [Interfaces Feature Guide for Security Devices](#)

Using the J-Web Ping Host Tool

Supported Platforms [SRX Series, vSRX](#)

You can ping a host to verify that the host can be reached over the network. The output is useful for diagnosing host and network connectivity problems. The device sends a series of ICMP echo (ping) requests to a specified host and receives ICMP echo responses.

Alternatively, you can use the CLI **ping** command. (See [“Using the ping Command” on page 550](#).)

To use the ping host tool:

1. Select **Troubleshoot>Ping Host** from the task bar.
2. Next to Advanced options, click the expand icon.
3. Enter information into the Ping Host page (see [Table 117 on page 553](#)).

Table 117: J-Web Ping Host Field Summary

Field	Function	Your Action
Remote Host	Identifies the host to ping. This is the only required field.	Type the hostname or IP address of the host to ping.
Advanced Options		
Don't Resolve Addresses	Determines whether to display hostnames of the hops along the path.	<ul style="list-style-type: none"> • Suppress the display of the hop hostnames by selecting the check box. • Display the hop hostnames by clearing the check box.
Interface	Specifies the interface on which the ping requests are sent.	Select the interface on which ping requests are sent from the list. If you select any , the ping requests are sent on all interfaces.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list.
Don't Fragment	Specifies the Don't Fragment (DF) bit in the IP header of the ping request packet.	<ul style="list-style-type: none"> • Set the DF bit by selecting the check box. • Clear the DF bit by clearing the check box.
Record Route	Sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed in the main pane.	<ul style="list-style-type: none"> • Record and display the path of the packet by selecting the check box. • Suppress the recording and display of the path of the packet by clearing the check box.
Type-of-Service	Specifies the type-of-service (TOS) value in the IP header of the ping request packet.	Select the decimal value of the TOS field from the list.
Routing Instance	Names the routing instance for the ping attempt.	Select the routing instance name from the list.
Interval	Specifies the interval, in seconds, between the transmission of each ping request.	Select the interval from the list.
Packet Size	Specifies the size of the ping request packet.	Type the size, in bytes, of the packet. The size can be from 0 through 65,468. The device adds 8 bytes of ICMP header to the size.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address.
Time-to-Live	Specifies the time-to-live (TTL) hop count for the ping request packet.	Select the TTL from the list.

Table 117: J-Web Ping Host Field Summary (*continued*)

Field	Function	Your Action
Bypass Routing	<p>Determines whether ping requests are routed by means of the routing table.</p> <p>If the routing table is not used, ping requests are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, ping responses are not sent.</p>	<ul style="list-style-type: none"> Bypass the routing table and send the ping requests to hosts on the specified interface only by selecting the check box. Route the ping requests using the routing table by clearing the check box.

4. Click **Start**.

The results of the ping operation appear in the main pane. If no options are specified, each ping response is in the following format:

bytes bytes from ip-address: icmp_seq=number ttl=number time=time

5. You can stop the ping operation before it is complete by clicking **OK**.**Related Documentation**

- [Diagnostic Tools Overview on page 8](#)
- [Understanding Ping MPLS on page 549](#)
- [J-Web Ping Host Results and Output Summary on page 554](#)
- [Using the J-Web Traceroute Tool on page 540](#)
- [Using the J-Web Ping MPLS Tool on page 555](#)
- [Using the J-Web Packet Capture Tool on page 586](#)

J-Web Ping Host Results and Output Summary

Supported Platforms [SRX Series, vSRX](#)

[Table 118 on page 554](#) summarizes the output in the ping host display.

Table 118: Ping Host Results and Output

Ping Host Result	Description
<i>bytes bytes from ip-address</i>	<ul style="list-style-type: none"> bytes—Size of ping response packet, which is equal to the value you entered in the Packet Size box, plus 8. ip-address—IP address of destination host that sent the ping response packet.
<i>icmp_seq=0</i> <i>icmp_seq=number</i>	number —Sequence Number field of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
<i>ttl=number</i>	number —Time-to-live hop-count value of the ping response packet.
<i>number packets transmitted</i>	number —Number of ping requests (probes) sent to host.

Table 118: Ping Host Results and Output (*continued*)

Ping Host Result	Description
<i>percentage packet loss</i>	<i>percentage</i> —Number of ping responses divided by the number of ping requests, specified as a percentage.
<i>round-trip min/avg/max/stddev = min-time/avg-time/max-time/std-dev ms</i>	<ul style="list-style-type: none"> <i>min-time</i>—Minimum round-trip time (see <i>time=time</i> field in this table). <i>avg-time</i>—Average round-trip time. <i>max-time</i>—Maximum round-trip time. <i>std-dev</i>—Standard deviation of the round-trip times.

If the device does not receive ping responses from the destination host (the output shows a packet loss of 100 percent), one of the following explanations might apply:

- The host is not operational.
- There are network connectivity problems between the device and the host.
- The host might be configured to ignore ICMP echo requests.
- The host might be configured with a firewall filter that blocks ICMP echo requests or ICMP echo responses.
- The size of the ICMP echo request packet exceeds the MTU of a host along the path.
- The value you selected in the Time-to-Live box was less than the number of hops in the path to the host, in which case the host might reply with an ICMP error message.

Related Documentation

- [Diagnostic Tools Overview on page 8](#)
- [Understanding Ping MPLS on page 549](#)
- [Using the J-Web Ping Host Tool on page 552](#)
- [Interfaces Feature Guide for Security Devices](#)

Using the J-Web Ping MPLS Tool

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

Before using the ping MPLS feature, make sure that the receiving interface on the VPN or LSP remote endpoint has MPLS enabled, and that the loopback interface on the outbound node is configured as **127.0.0.1**. The source address for MPLS probes must be a valid address on the device.

To use the ping MPLS tool:

1. Select **Troubleshoot>Ping MPLS** from the task bar.
2. Next to the ping MPLS option you want to use, click the expand icon.
3. Enter information into the Ping MPLS page (see [Table 119 on page 556](#)).

Table 119: J-Web Ping MPLS Field Summary

Field	Function	Your Action
Ping RSVP-signaled LSP		
LSP Name	Identifies the LSP to ping.	Type the name of the LSP to ping.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list. The default is 5 requests.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
Ping LDP-signaled LSP		
FEC Prefix	Identifies the LSP to ping.	Type the forwarding equivalence class (FEC) prefix and length of the LSP to ping.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list. The default is 5 requests.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
Ping LSP to Layer 3 VPN prefix		
Layer 3 VPN Name	Identifies the Layer 3 VPN to ping.	Type the name of the VPN to ping.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list. The default is 5 requests.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
VPN Prefix	Identifies the IP address prefix and length of the Layer 3 VPN to ping.	Type the IP address prefix and length of the VPN to ping.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.
Locate LSP using interface name		
Interface	Specifies the interface on which the ping requests are sent.	Select the device interface on which ping requests are sent from the list. If you select any , the ping requests are sent on all interfaces.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.

Table 119: J-Web Ping MPLS Field Summary (*continued*)

Field	Function	Your Action
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list. The default is 5 requests.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
Instance to which this connection belongs		
Layer 2VPN Name	Identifies the Layer 2 VPN to ping.	Type the name of the VPN to ping.
Remote Site Identifier	Specifies the remote site identifier of the Layer 2 VPN to ping.	Type the remote site identifier for the VPN.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.
Local Site Identifier	Specifies the local site identifier of the Layer 2 VPN to ping.	Type the local site identifier for the VPN.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list. The default is 5 requests.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
Locate LSP from interface name		
Interface	Specifies the interface on which the ping requests are sent.	Select the device interface on which ping requests are sent from the list. If you select any , the ping requests are sent on all interfaces.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list. The default is 5 requests.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
Locate LSP from virtual circuit information		
Remote Neighbor	Identifies the remote neighbor (PE device) within the virtual circuit to ping.	Type the IP address of the remote neighbor within the virtual circuit.
Circuit Identifier	Specifies the virtual circuit identifier for the Layer 2 circuit to ping.	Type the virtual circuit identifier for the Layer 2 circuit.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.

Table 119: J-Web Ping MPLS Field Summary (*continued*)

Field	Function	Your Action
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.
Ping end point of LSP		
VPN Prefix	Identifies the LSP endpoint to ping.	Type either the LDP FEC prefix and length or the RSVP LSP endpoint address for the LSP to ping.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address—a valid address configured on a device interface.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list.
Detailed Output	Requests the display of extensive rather than brief ping output.	Select the check box to display detailed output.

4. Click **Start**.

5. You can stop the ping operation before it is complete by clicking **OK**.

Related Documentation

- [Diagnostic Tools Overview on page 8](#)
- [Understanding Ping MPLS on page 549](#)
- [J-Web Ping MPLS Results and Output Summary on page 558](#)
- [Using the J-Web Traceroute Tool on page 540](#)
- [Using the J-Web Ping Host Tool on page 552](#)
- [Using the J-Web Packet Capture Tool on page 586](#)

J-Web Ping MPLS Results and Output Summary

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

[Table 120 on page 558](#) summarizes the output in the ping MPLS display.

Table 120: J-Web Ping MPLS Results and Output Summary

Field	Description
Exclamation point (!)	Echo reply was received.

Table 120: J-Web Ping MPLS Results and Output Summary (*continued*)

Field	Description
Period (.)	Echo reply was not received within the timeout period.
x	Echo reply was received with an error code. Errored packets are not counted in the received packets count and are accounted for separately.
<i>number packets transmitted</i>	<i>number</i> —Number of ping requests (probes) sent to a host.
<i>number packets received</i>	<i>number</i> —Number of ping responses received from a host.
<i>percentage packet loss</i>	<i>percentage</i> —Number of ping responses divided by the number of ping requests, specified as a percentage.
time	For Layer 2 circuits only, the number of milliseconds required for the ping packet to reach the destination. This value is approximate, because the packet has to reach the Routing Engine.

If the device does not receive ping responses from the destination host (the output shows a packet loss of 100 percent), one of the following explanations might apply:

- The host is not operational.
- There are network connectivity problems between the device and the host.
- The host might be configured to ignore echo requests.
- The host might be configured with a firewall filter that blocks echo requests or echo responses.
- The size of the echo request packet exceeds the MTU of a host along the path.
- The outbound node at the remote endpoint is not configured to handle MPLS packets.
- The remote endpoint's loopback address is not configured to 127.0.0.1.

Related Documentation

- [Diagnostic Tools Overview on page 8](#)
- [Understanding Ping MPLS on page 549](#)
- [Using the J-Web Ping MPLS Tool on page 555](#)
- *Interfaces Feature Guide for Security Devices*

Pinging Layer 2 Circuits

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

Enter the **ping mpls l2circuit** command with the following syntax:

```
user@host> ping mpls l2circuit (interface interface-name | virtual-circuit neighbor
    prefix-name virtual-circuit-id) <exp forwarding-class> <count number>
    <source source-address> <detail>
```

Table 121 on page 560 describes the **ping mpls l2circuit** command options.

Table 121: CLI ping mpls l2circuit Command Options

Option	Description
l2circuit interface <i>interface-name</i>	Sends ping requests out the specified interface configured for the Layer 2 circuit on the outbound PE device.
l2circuit virtual-circuit neighbor <i>prefix-name</i> <i>virtual-circuit-id</i>	Pings on a combination of the IPv4 prefix and the virtual circuit identifier on the outbound PE device, testing the integrity of the Layer 2 circuit between the inbound and outbound PE devices.
exp <i>forwarding-class</i>	(Optional) Specifies the value of the forwarding class to be used in the MPLS ping packets.
count <i>number</i>	(Optional) Limits the number of ping requests to send. Specify a count from 0 through 1,000,000. The default value is 5. If you do not specify a count, ping requests are continuously sent until you press Ctrl-C.
source <i>source-address</i>	(Optional) Uses the source address that you specify, in the ping request packet.
detail	(Optional) Displays detailed output about the echo requests sent and received. Detailed output includes the MPLS labels used for each request and the return codes for each request.

To quit the **ping mpls l2circuit** command, press Ctrl-C.

The following is sample output from a **ping mpls l2circuit** command:

```
user@host> ping mpls l2circuit interface fe-1/0/0.0
```

```
Request for seq 1, to interface 69, labels <100000, 100208>
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

The fields in the display are the same as those displayed by the J-Web ping MPLS diagnostic tool.

Related Documentation

- [Using the ping Command on page 550](#)
- [Understanding Ping MPLS on page 549](#)
- [Pinging Layer 2 VPNs on page 560](#)
- [Pinging Layer 3 VPNs on page 562](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 563](#)
- [Using the J-Web Ping Host Tool on page 552](#)

Pinging Layer 2 VPNs

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

Enter the **ping mpls l2vpn** command with the following syntax:

```

user@host> ping mpls l2vpn interface interface-name | instance l2vpn-instance-name
local-site-id local-site-id-number remote-site-id remote-site-id-number
<bottom-label-ttl> <exp forwarding-class> <count number> <source source-address>
<detail>

```

Table 122 on page 561 describes the **ping mpls l2vpn** command options.

Table 122: CLI ping mpls l2vpn Command Options

Option	Description
l2vpn interface <i>interface-name</i>	Sends ping requests out the specified interface configured for the Layer 2 VPN on the outbound (egress) PE device.
l2vpn instance <i>l2vpn-instance-name</i> <i>local-site-id</i> <i>local-site-id-number</i> <i>remote-site-id</i> <i>remote-site-id-number</i>	Pings on a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier, testing the integrity of the Layer 2 VPN circuit (specified by the identifiers) between the inbound (ingress) and outbound PE devices.
bottom-label-ttl	(Optional) Displays the time-to-live (TTL) value for the bottom label in the MPLS label stack.
exp forwarding-class	(Optional) Specifies the value of the forwarding class to be used in the MPLS ping packets.
countnumber	(Optional) Limits the number of ping requests to send. Specify a count from 0 through 1,000,000. The default value is 5. If you do not specify a count, ping requests are continuously sent until you press Ctrl-C.
source source-address	(Optional) Uses the source address that you specify, in the ping request packet.
detail	(Optional) Displays detailed output about the echo requests sent and received. Detailed output includes the MPLS labels used for each request and the return codes for each request.

To quit the **ping mpls l2vpn** command, press Ctrl-C.

The following is sample output from a **ping mpls l2vpn** command:

```

user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2 detail

Request for seq 1, to interface 68, labels <800001, 100176>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <800001, 100176>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <800001, 100176>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <800001, 100176>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <800001, 100176>
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss

```

The fields in the display are the same as those displayed by the J-Web ping MPLS diagnostic tool.

Related Documentation

- [Using the ping Command on page 550](#)
- [Understanding Ping MPLS on page 549](#)
- [Pinging Layer 2 Circuits on page 559](#)
- [Pinging Layer 3 VPNs on page 562](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 563](#)
- [Using the J-Web Ping Host Tool on page 552](#)

Pinging Layer 3 VPNs

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

Enter the **ping mpls l3vpn** command with the following syntax:

```
user@host> ping mpls l3vpn prefix prefix-name <l3vpn-name> <bottom-label-ttl>
<exp forwarding-class> <count number> <source source-address> <detail>
```

[Table 123 on page 562](#) describes the **ping mpls l3vpn** command options.

Table 123: CLI ping mpls l3vpn Command Options

Option	Description
l3vpn prefix <i>prefix-name</i>	Pings the remote host specified by the prefix to verify that the prefix is present in the PE device's VPN routing and forwarding (VRF) table. This option does not test the connectivity between a PE device and a CE device.
<i>l3vpn-name</i>	(Optional) Layer 3 VPN name.
bottom-label-ttl	(Optional) Displays the time-to-live (TTL) value for the bottom label in the MPLS label stack.
exp <i>forwarding-class</i>	(Optional) Specifies the value of the forwarding class to be used in the MPLS ping packets.
count <i>number</i>	(Optional) Limits the number of ping requests to send. Specify a count from 0 through 1,000,000. The default value is 5. If you do not specify a count, ping requests are continuously sent until you press Ctrl-C.
source <i>source-address</i>	(Optional) Uses the source address that you specify, in the ping request packet.
detail	(Optional) Displays detailed output about the echo requests sent and received. Detailed output includes the MPLS labels used for each request and the return codes for each request.

To quit the **ping mpls l3vpn** command, press Ctrl-C.

The following is sample output from a **ping mpls l3vpn** command:

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
```

```
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

The fields in the display are the same as those displayed by the J-Web ping MPLS diagnostic tool.

Related Documentation

- [Using the ping Command on page 550](#)
- [Understanding Ping MPLS on page 549](#)
- [Pinging Layer 2 Circuits on page 559](#)
- [Pinging Layer 2 VPNs on page 560](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 563](#)
- [Using the J-Web Ping Host Tool on page 552](#)

Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

Enter the **ping mpls** command with the following syntax:

```
user@host> ping mpls (ldp fec | lsp-end-point prefix-name | rsvp lsp-name)
<exp forwarding-class> <count number> <source source-address> <detail>
```

[Table 124 on page 563](#) describes the **ping mpls** command options.

Table 124: CLI ping mpls ldp and ping mpls lsp-end-point Command Options

Option	Description
ldp fec	Pings an LDP-signaled LSP identified by the forwarding equivalence class (FEC) prefix and length.
lsp-end-point prefix-name	Pings an LSP endpoint using either an LDP FEC or a RSVP LSP endpoint address.
rsvp lsp-name	Pings an RSVP-signaled LSP identified by the specified LSP name.
exp forwarding-class	(Optional) Specifies the value of the forwarding class to be used in the MPLS ping packets.
countnumber	(Optional) Limits the number of ping requests to send. Specify a count from 0 through 1,000,000. The default value is 5. If you do not specify a count, ping requests are continuously sent until you press Ctrl-C.
source source-address	(Optional) Uses the source address that you specify, in the ping request packet.
detail	(Optional) Displays detailed output about the echo requests sent and received. Detailed output includes the MPLS labels used for each request and the return codes for each request.

To quit the **ping mpls** command, press Ctrl-C.

The following is sample output from a **ping mpls** command:

```
user@host> ping mpls rsvp count 5

!!xxx
--- 1sping statistics ---
5 packets transmitted, 2 packets received, 60% packet loss
3 packets received with error status, not counted as received.
```

The fields in the display are the same as those displayed by the J-Web ping MPLS diagnostic tool.

**Related
Documentation**

- [Using the ping Command on page 550](#)
- [Understanding Ping MPLS on page 549](#)
- [Pinging Layer 2 Circuits on page 559](#)
- [Pinging Layer 2 VPNs on page 560](#)
- [Pinging Layer 3 VPNs on page 562](#)
- [Using the J-Web Ping Host Tool on page 552](#)

CHAPTER 33

Using Packet Capture to Analyze Network Traffic

- [Packet Capture Overview on page 565](#)
- [Example: Enabling Packet Capture on a Device on page 568](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
- [Example: Configuring Packet Capture for Datapath Debugging on page 575](#)
- [Disabling Packet Capture on page 579](#)
- [Deleting Packet Capture Files on page 579](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 580](#)
- [Displaying Packet Headers on page 581](#)
- [Using the J-Web Packet Capture Tool on page 586](#)
- [J-Web Packet Capture Results and Output Summary on page 589](#)

Packet Capture Overview

Supported Platforms [SRX Series, vSRX](#)

Packet capture is a tool that helps you to analyze network traffic and troubleshoot network problems. The packet capture tool captures real-time data packets traveling over the network for monitoring and logging.



NOTE: Packet capture is supported on physical interfaces, reth interfaces, and tunnel interfaces, such as gr, ip, st0, and lsq-/ls.

Packets are captured as binary data, without modification. You can read the packet information offline with a packet analyzer such as Ethereal or tcpdump. If you need to quickly capture packets destined for, or originating from, the Routing Engine and analyze them online, you can use the J-Web packet capture diagnostic tool.



NOTE: The packet capture tool does not support IPv6 packet capture.

You can use either the J-Web configuration editor or CLI configuration editor to configure packet capture.

Network administrators and security engineers use packet capture to perform the following tasks:

- Monitor network traffic and analyze traffic patterns.
- Identify and troubleshoot network problems.
- Detect security breaches in the network, such as unauthorized intrusions, spyware activity, or ping scans.

Packet capture operates like traffic sampling on the device, except that it captures entire packets including the Layer 2 header and saves the contents to a file in libpcap format. Packet capture also captures IP fragments. You cannot enable packet capture and traffic sampling on the device at the same time. Unlike traffic sampling, there are no tracing operations for packet capture.



NOTE: You can enable packet capture and port mirroring simultaneously on a device.

This section contains the following topics:

- [Packet Capture on Device Interfaces on page 566](#)
- [Firewall Filters for Packet Capture on page 567](#)
- [Packet Capture Files on page 567](#)
- [Analysis of Packet Capture Files on page 567](#)

Packet Capture on Device Interfaces

Packet capture is supported on the T1, T3, E1, E3, serial, Fast Ethernet, ADSL, G.SHDSL, PPPoE, and ISDN interfaces.

To capture packets on an ISDN interface, configure packet capture on the dialer interface. To capture packets on a PPPoE interface, configure packet capture on the PPPoE logical interface.

Packet capture supports PPP, Cisco HDLC, Frame Relay, and other ATM encapsulations. Packet capture also supports Multilink PPP (MLPPP), Multilink Frame Relay end-to-end (MLFR), and Multilink Frame Relay UNI/NNI (MFR) encapsulations.

You can capture all IPv4 packets flowing on an interface in the inbound or outbound direction. However, on traffic that bypasses the flow software module (protocol packets such as ARP, OSPF, and PIM), packets generated by the Routing Engine are not captured unless you have configured and applied a firewall filter on the interface in the outbound direction.

Tunnel interfaces can support packet capture in the outbound direction only.

Use the J-Web configuration editor or CLI configuration editor to specify the maximum packet size, the filename to be used for storing the captured packets, the maximum file size, the maximum number of packet capture files, and the file permissions.



NOTE: For packets captured on T1, T3, E1, E3, serial, and ISDN interfaces in the outbound (egress) direction, the size of the packet captured might be 1 byte less than the maximum packet size configured because of the packet loss priority (PLP) bit.

To modify encapsulation on an interface that has packet capture configured, you must first disable packet capture.

Firewall Filters for Packet Capture

When you enable packet capture on a device, all packets flowing in the direction specified in packet capture configuration (inbound, outbound, or both) are captured and stored. Configuring an interface to capture all packets might degrade the performance of the device. You can control the number of packets captured on an interface with firewall filters and specify various criteria to capture packets for specific traffic flows.

You must also configure and apply appropriate firewall filters on the interface if you need to capture packets generated by the host device, because interface sampling does not capture packets originating from the host device.

Packet Capture Files

When packet capture is enabled on an interface, the entire packet including the Layer 2 header is captured and stored in a file. You can specify the maximum size of the packet to be captured, up to 1500 bytes. Packet capture creates one file for each physical interface. You can specify the target filename, the maximum size of the file, and the maximum number of files.

File creation and storage take place in the following way. Suppose you name the packet capture file **pcap-file**. Packet capture creates multiple files (one per physical interface), suffixing each file with the name of the physical interface; for example, **pcap-file.fe-0.0.1** for the Fast Ethernet interface **fe-0.0.1**. When the file named **pcap-file.fe-0.0.1** reaches the maximum size, the file is renamed **pcap-file.fe-0.0.1.0**. When the file named **pcap-file.fe-0.0.1** reaches the maximum size again, the file named **pcap-file.fe-0.0.1.0** is renamed **pcap-file.fe-0.0.1.1** and **pcap-file.fe-0.0.1** is renamed **pcap-file.fe-0.0.1.0**. This process continues until the maximum number of files is exceeded and the oldest file is overwritten. The **pcap-file.fe-0.0.1** file is always the latest file.

Packet capture files are not removed even after you disable packet capture on an interface.

Analysis of Packet Capture Files

Packet capture files are stored in libpcap format in the **/var/tmp** directory. You can specify user or administrator privileges for the files.

Packet capture files can be opened and analyzed offline with tcpdump or any packet analyzer that recognizes the libpcap format. You can also use FTP or the Session Control Protocol (SCP) to transfer the packet capture files to an external device.



NOTE: Disable packet capture before opening the file for analysis or transferring the file to an external device with FTP or SCP. Disabling packet capture ensures that the internal file buffer is flushed and all the captured packets are written to the file.

**Related
Documentation**

- [Example: Enabling Packet Capture on a Device on page 568](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
- [Using the J-Web Packet Capture Tool on page 586](#)

Example: Enabling Packet Capture on a Device

Supported Platforms [SRX Series, vSRX](#)

This example shows how to enable packet capture on a device, allowing you to analyze network traffic and troubleshoot network problems

- [Requirements on page 568](#)
- [Overview on page 568](#)
- [Configuration on page 569](#)
- [Verification on page 570](#)

Requirements

Before you begin:

- Establish basic connectivity.
- Configure network interfaces. See [Interfaces Feature Guide for Security Devices](#).

Overview

In this example, you set the maximum packet capture size in each file as 500 bytes. The range is from 68 through 1500, and the default is 68 bytes. You specify the target filename for the packet capture file as pcap-file. You then specify the maximum number of files to capture as 100. The range is from 2 through 10,000, and the default is 10 files. You set the maximum size of each file to 1024 bytes. The range is from 1,024 through 104,857,600, and the default is 512,000 bytes. Finally, you specify that all users have permission to read the packet capture files.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set forwarding-options packet-capture maximum-capture-size 500
set forwarding-options packet-capture file filename pcap-file files 100 size 1024
world-readable
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To enable packet capture on a device:

1. Set the maximum packet capture size.

```
[edit]
user@host# edit forwarding-options
user@host# set packet-capture maximum-capture-size 500
```

2. Specify the target filename.

```
[edit forwarding-options]
user@host# set packet-capture file filename pcap-file
```

3. Specify the maximum number of files to capture.

```
[edit forwarding-options]
user@host# set packet-capture file files 100
```

4. Specify the maximum size of each file.

```
[edit forwarding-options]
user@host# set packet-capture file size 1024
```

5. Specify that all users have permission to read the file.

```
[edit forwarding-options]
user@host# set packet-capture file world-readable
```

Results From configuration mode, confirm your configuration by entering the **show forwarding-options** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show forwarding-options
packet-capture {
  file filename pcap-file files 100 size 1k world-readable;
```

```
        maximum-capture-size 500;  
    }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying the Packet Capture Configuration on page 570](#)
- [Verifying Captured Packets on page 570](#)

Verifying the Packet Capture Configuration

Purpose Verify that the packet capture is configured on the device.

Action From configuration mode, enter the **show forwarding-options** command. Verify that the output shows the intended file configuration for capturing packets.

Verifying Captured Packets

Purpose Verify that the packet capture file is stored under the **/var/tmp** directory and the packets can be analyzed offline.

Action 1. Disable packet capture.

Using FTP, transfer a packet capture file (for example, **126b.fe-0.0.1**), to a server where you have installed packet analyzer tools (for example, **tools-server**).

a. From configuration mode, connect to **tools-server** using FTP.

```
[edit]  
user@host# run ftp tools-server  
Connected to tools-server.mydomain.net  
220 tools-server.mydomain.net FTP server (Version 6.00LS) ready  
Name (tools-server:user):remoteuser  
331 Password required for remoteuser.  
Password:  
230 User remoteuser logged in.  
Remote system type is UNIX.  
Using binary mode to transfer files.  
ftp>
```

b. Navigate to the directory where packet capture files are stored on the device.

```
ftp> lcd /var/tmp  
Local directory now /cf/var/tmp
```

c. Copy the packet capture file that you want to analyze to the server, for example **126b.fe-0.0.1**.

```
ftp> put 126b.fe-0.0.1
local: 126b.fe-0.0.1 remote: 126b.fe-0.0.1
200 PORT command successful.
150 Opening BINARY mode data connection for '126b.fe-0.0.1'.
100% 1476 00:00 ETA
226 Transfer complete.
1476 bytes sent in 0.01 seconds (142.42 KB/s)
```

d. Return to configuration mode.

```
ftp> bye
221 Goodbye.
[edit]
user@host#
```

2. Open the packet capture file on the server with tcpdump or any packet analyzer that supports libpcap format and review the output.

```
root@server% tcpdump -r 126b.fe-0.0.1 -xvvvv
```

```
01:12:36.279769 Out 0:5:85:c4:e3:d1 > 0:5:85:c8:f6:d1, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 64, id 33133, offset 0, flags [none], proto: ICMP (1),
length: 84) 14.1.1.1 > 15.1.1.1: ICMP echo request seq 0, length 64
    0005 85c8 f6d1 0005 85c4 e3d1 0800 4500
    0054 816d 0000 4001 da38 0e01 0101 0f01
    0101 0800 3c5a 981e 0000 8b5d 4543 51e6
    0100 aaaa aaaa aaaa aaaa aaaa aaaa aaaa
    aaaa aaaa 0000 0000 0000 0000 0000 0000
    0000 0000 0000 0000 0000 0000 0000 0000
    0000
01:12:36.279793 Out 0:5:85:c8:f6:d1 > 0:5:85:c4:e3:d1, ethertype IPv4 (0x0800),
length 98: (tos 0x0, ttl 63, id 41227, offset 0, flags [none], proto: ICMP (1),
length: 84) 15.1.1.1 > 14.1.1.1: ICMP echo reply seq 0, length 64
    0005 85c4 e3d1 0005 85c8 f6d1 0800 4500
    0054 a10b 0000 3f01 bb9a 0f01 0101 0e01
    0101 0000 445a 981e 0000 8b5d 4543 51e6
    0100 aaaa aaaa aaaa aaaa aaaa aaaa aaaa
    aaaa aaaa 0000 0000 0000 0000 0000 0000
    0000 0000 0000 0000 0000 0000 0000 0000
    0000
```

```
root@server%
```

Related Documentation

- [Packet Capture Overview on page 565](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
- [Disabling Packet Capture on page 579](#)
- [Deleting Packet Capture Files on page 579](#)
- [Disabling Packet Capture on page 579](#)

Example: Configuring Packet Capture on an Interface

Supported Platforms SRX Series, vSRX

This example shows how to configure packet capture on an interface to analyze traffic.

- [Requirements on page 572](#)
- [Overview on page 572](#)
- [Configuration on page 572](#)
- [Verification on page 573](#)

Requirements

Before you begin:

- Establish basic connectivity.
- Configure network interfaces. See [Interfaces Feature Guide for Security Devices](#).

Overview

In this example, you create an interface called fe-0/0/1. You then configure the direction of the traffic for which you are enabling packet capture on the logical interface as inbound and outbound.



NOTE: On traffic that bypasses the flow software module (protocol packets such as ARP, OSPF, and PIM), packets generated by the Routing Engine are not captured unless you have configured and applied a firewall filter on the interface in the output direction.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
edit interfaces fe-0/0/1
set unit 0 family inet sampling input output
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure packet capture on an interface:

1. Create an interface.
 [edit]

```
user@host# edit interfaces fe-0/0/1
```

2. Configure the direction of the traffic.

```
[edit interfaces fe-0/0/1]
user@host# set unit 0 family inet sampling input output
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

Verifying the Packet Capture Configuration

Purpose Confirm that the configuration is working properly.

Verify that packet capture is configured on the interface.

Action From configuration mode, enter the **show interfaces fe-0/0/1** command.

Related Documentation

- [Packet Capture Overview on page 565](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 580](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
- [Example: Enabling Packet Capture on a Device on page 568](#)
- [Deleting Packet Capture Files on page 579](#)
- [Disabling Packet Capture on page 579](#)

Example: Configuring a Firewall Filter for Packet Capture

Supported Platforms [SRX Series, vSRX](#)

This example shows how to configure a firewall filter for packet capture and apply it to a logical interface.

- [Requirements on page 573](#)
- [Overview on page 574](#)
- [Configuration on page 574](#)
- [Verification on page 575](#)

Requirements

Before you begin:

- Establish basic connectivity.
- Configure network interfaces. See [Interfaces Feature Guide for Security Devices](#).

Overview

In this example, you set a firewall filter called `dest-all` and a term name called `dest-term` to capture packets from a specific destination address, which is `192.168.1.1/32`. You define the match condition to accept the sampled packets. Finally, you apply the `dest-all` filter to all of the outgoing packets on interface `fe-0/0/1`.



NOTE: If you apply a firewall filter on the loopback interface, it affects all traffic to and from the Routing Engine. If the firewall filter has a `sample` action, packets to and from the Routing Engine are sampled. If packet capture is enabled, then packets to and from the Routing Engine are captured in the files created for the input and output interfaces.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set firewall filter dest-all term dest-term from destination-address 192.168.1.1/32
set firewall filter dest-all term dest-term then sample accept
edit interfaces
set interfaces fe-0/0/1 unit 0 family inet filter output dest-all
```

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a firewall filter for packet capture and apply it to a logical interface:

1. Specify the firewall filter and its destination address.

```
[edit]
user@host# edit firewall
user@host# set filter dest-all term dest-term from destination-address 192.168.1.1/32
```
2. Define the match condition and its action.

```
[edit firewall]
user@host# set filter dest-all term dest-term then sample accept
```
3. Apply the filter to all the outgoing packets.

```
[edit interfaces]
user@host# set interfaces fe-0/0/1 unit 0 family inet filter output dest-all
```


Results From configuration mode, confirm your configuration by entering the **show firewall filter dest-all** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show firewall filter dest-all
term dest-term {
  from {
    destination-address 192.168.1.1/32;
  }
  then {
    sample;
    accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Verifying the Firewall Filter for Packet Capture Configuration

Purpose Confirm that the configuration is working properly.

Verify that the firewall filter for packet capture is configured.

Action From configuration mode, enter the **show firewall filter dest-all** command. Verify that the output shows the intended configuration of the firewall filter for capturing packets sent to the destination address.

Related Documentation

- [Packet Capture Overview on page 565](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)
- [Example: Enabling Packet Capture on a Device on page 568](#)
- [Deleting Packet Capture Files on page 579](#)
- [Disabling Packet Capture on page 579](#)

Example: Configuring Packet Capture for Datapath Debugging

Supported Platforms [SRX Series, vSRX](#)

This example shows how to configure packet capture to monitor traffic that passes through the device. Packet capture then dumps the packets into a PCAP file format that can be later examined by the tcpdump utility.

- [Requirements on page 576](#)
- [Overview on page 576](#)

- [Configuration on page 576](#)
- [Verification on page 578](#)

Requirements

Before you begin, see “[Debugging the Data Path \(CLI Procedure\)](#)” on page 530.

Overview

A filter is defined to filter traffic; then an action profile is applied to the filtered traffic. The action profile specifies a variety of actions on the processing unit. One of the supported actions is packet dump, which sends the packet to the Routing Engine and stores it in proprietary form to be read using the **show security datapath-debug capture** command.



NOTE: Datapath debugging is not supported on SRX1500 devices.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set security datapath-debug capture-file my-capture
set security datapath-debug capture-file format pcap
set security datapath-debug capture-file size 1m
set security datapath-debug capture-file files 5
set security datapath-debug maximum-capture-size 400
set security datapath-debug action-profile do-capture event np-ingress packet-dump
set security datapath-debug packet-filter my-filter action-profile do-capture
set security datapath-debug packet-filter my-filter source-prefix 1.2.3.4/32
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure packet capture:

1. Edit the security datapath-debug option for the multiple processing units along the packet-processing path:

```
[edit]
user@host# edit security datapath-debug
```
2. Enable the capture file, the file format, the file size, and the number of files. Size number limits the size of the capture file. After the limit size is reached, if the file number is specified, then the capture file will be rotated to filename *x*, where *x* is auto-incremented until it reaches the specified index and then returns to zero. If no

files index is specified, the packets will be discarded after the size limit is reached. The default size is 512 kilobytes.

```
[edit security datapath-debug]
user@host# set capture-file my-capture format pcap size 1m files 5
[edit security datapath-debug]
user@host# set maximum-capture-size 400
```

3. Enable action profile and set the event. Set the action profile as do-capture and the event type as np-ingress:

```
[edit security datapath-debug]
user@host# edit action-profile do-capture
[edit security datapath-debug action-profile do-capture]
user@host# edit event np-ingress
```

4. Enable packet dump for the action profile:

```
[edit security datapath-debug action-profile do-capture event np-ingress]
user@host# set packet-dump
```

5. Enable packet filter, action, and filter options. The packet filter is set to my-filter, the action profile is set to do-capture, and filter option is set to source-prefix 1.2.3.4/32.

```
[edit security datapath-debug]
user@host# set security datapath-debug packet-filter my-filter action-profile
do-capture

[edit security datapath-debug]
user@host# set security datapath-debug packet-filter my-filter source-prefix
1.2.3.4/32
```

Results From configuration mode, confirm your configuration by entering the **show security datapath-debug** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it. The following is **show security datapath-debug** output from the **show security datapath-debug** command:

```
security {
  datapath-debug {
    capture-file {
      my-capture
      format pcap
      size 1m
      files 5;
    }
  }
  maximum-capture-size 100;
  action-profile do-capture {
    event np-ingress {
      packet-dump
    }
  }
}
```

```
packet-filter my-filter {  
  source-prefix 1.2.3.4/32  
  action-profile do-capture  
}  
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying Packet Capture on page 578](#)
- [Verifying Data Path Debugging Capture on page 578](#)
- [Verifying Data Path Debugging Counter on page 578](#)

Verifying Packet Capture

Purpose Verify if the packet capture is working.

Action From operational mode, enter the **request security datapath-debug capture start** command to start packet capture and enter the **request security datapath-debug capture stop** command to stop packet capture.

To view the results, from CLI operational mode, access the local UNIX shell and navigate to the directory `/var/log/my-capture`. The result can be read by using the `tcpdump` utility.

Verifying Data Path Debugging Capture

Purpose Verify the details of data path debugging capture file.

Action From operational mode, enter the **show security datapath-debug capture** command.

```
user@host>show security datapath-debug capture
```



WARNING: When you are done troubleshooting, make sure to remove or deactivate all the traceoptions configurations (not limited to flow traceoptions) and the complete security datapath-debug configuration stanza. If any debugging configurations remain active, they will continue to use the device's CPU and memory resources.

Verifying Data Path Debugging Counter

Purpose Verify the details of the data path debugging counter.

Action From operational mode, enter the [show security datapath-debug counter](#) command.

- Related Documentation**
- [Packet Capture Overview on page 565](#)
 - [Understanding Data Path Debugging for SRX Series Devices on page 529](#)
 - [Debugging the Data Path \(CLI Procedure\) on page 530](#)

Disabling Packet Capture

Supported Platforms [SRX Series, vSRX](#)

You must disable packet capture before opening the packet capture file for analysis or transferring the file to an external device. Disabling packet capture ensures that the internal file buffer is flushed and all the captured packets are written to the file.

To disable packet capture, enter from configuration mode:

```
[edit forwarding-options]
user@host# set packet-capture disable
```

If you are done configuring the device, enter **commit** from configuration mode.

- Related Documentation**
- [Packet Capture Overview on page 565](#)
 - [Example: Configuring Packet Capture on an Interface on page 572](#)
 - [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
 - [Example: Enabling Packet Capture on a Device on page 568](#)
 - [Deleting Packet Capture Files on page 579](#)

Deleting Packet Capture Files

Supported Platforms [SRX Series, vSRX](#)

Deleting packet capture files from the /var/tmp directory only temporarily removes the packet capture files. Packet capture files for the interface are automatically created again the next time a packet capture configuration change is committed or as part of a packet capture file rotation.

To delete a packet capture file:

1. Disable packet capture (see [“Disabling Packet Capture” on page 579](#)).
2. Delete the packet capture file for the interface.
 - a. From operational mode, access the local UNIX shell.

```
user@host> start shell
%
```

- b. Navigate to the directory where packet capture files are stored.

```
% cd /var/tmp
%
```

- c. Delete the packet capture file for the interface; for example **pcap-file.fe.0.0.0**.

```
% rm pcap-file.fe.0.0.0
%
```

- d. Return to operational mode.

```
% exit
user@host>
```

3. Reenable packet capture (see [“Example: Enabling Packet Capture on a Device” on page 568](#)).

4. If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- [Packet Capture Overview on page 565](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
- [Example: Enabling Packet Capture on a Device on page 568](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 580](#)
- [Disabling Packet Capture on page 579](#)

Changing Encapsulation on Interfaces with Packet Capture Configured

Supported Platforms [SRX Series, vSRX](#)

Before modifying the encapsulation on a device interface that is configured for packet capture, you must disable packet capture and rename the latest packet capture file. Otherwise, packet capture saves the packets with different encapsulations in the same packet capture file. Packet files containing packets with different encapsulations are not useful, because packet analyzer tools like tcpdump cannot analyze such files.

After modifying the encapsulation, you can safely reenabling packet capture on the device.

To change the encapsulation on interfaces with packet capture configured:

1. Disable packet capture (see [“Disabling Packet Capture” on page 579](#)).
2. Enter **commit** from configuration mode.

3. Rename the latest packet capture file on which you are changing the encapsulation with the **.chdsl** extension.

- a. From operational mode, access the local UNIX shell.

```
user@host> start shell
%
```

- b. Navigate to the directory where packet capture files are stored.

```
% cd /var/tmp
%
```

- c. Rename the latest packet capture file for the interface on which you are changing the encapsulation; for example **fe.0.0.0**.

```
% mv pcap-file.fe.0.0.0 pcap-file.fe.0.0.0.chdsl
%
```

- d. Return to operational mode.

```
% exit
user@host>
```

4. Change the encapsulation on the interface using the J-Web user interface or CLI configuration editor.

5. If you are done configuring the device, enter **commit** from configuration mode.

6. Reenable packet capture (see [“Example: Enabling Packet Capture on a Device” on page 568](#)).

7. If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- [Packet Capture Overview on page 565](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 573](#)
- [Example: Enabling Packet Capture on a Device on page 568](#)

Displaying Packet Headers

Supported Platforms [SRX Series, vSRX](#)

Enter the **monitor traffic** command to display packet headers transmitted through network interfaces with the following syntax:



NOTE: Using the `monitor traffic` command can degrade system performance. We recommend that you use filtering options—such as `count` and `matching`—to minimize the impact to packet throughput on the system.

```
user@host> monitor traffic <absolute-sequence> <count number>
<interface interface-name> <layer2-headers> <matching "expression">
<no-domain-names> <no-promiscuous> <no-resolve> <no-timestamp> <print-ascii>
<print-hex> <size bytes> <brief | detail | extensive>
```

Table 125 on page 582 describes the `monitor traffic` command options.

Table 125: CLI monitor traffic Command Options

Option	Description
<code>absolute-sequence</code>	(Optional) Displays the absolute TCP sequence numbers.
<code>count number</code>	(Optional) Displays the specified number of packet headers. Specify a value from 0 through 100,000. The command quits and exits to the command prompt after this number is reached.
<code>interface interface-name</code>	(Optional) Displays packet headers for traffic on the specified interface. If an interface is not specified, the lowest numbered interface is monitored.
<code>layer2-headers</code>	(Optional) Displays the link-layer packet header on each line.
<code>matching "expression"</code>	(Optional) Displays packet headers that match an expression enclosed in quotation marks (" "). Table 126 on page 583 through Table 128 on page 585 list match conditions, logical operators, and arithmetic, binary, and relational operators you can use in the expression.
<code>no-domain-names</code>	(Optional) Suppresses the display of the domain name portion of the hostname.
<code>no-promiscuous</code>	(Optional) Specifies <i>not</i> to place the monitored interface in promiscuous mode. In promiscuous mode, the interface reads every packet that reaches it. In nonpromiscuous mode, the interface reads only the packets addressed to it.
<code>no-resolve</code>	(Optional) Suppresses the display of hostnames.
<code>no-timestamp</code>	(Optional) Suppresses the display of packet header timestamps.
<code>print-ascii</code>	(Optional) Displays each packet header in ASCII format.
<code>print-hex</code>	(Optional) Displays each packet header, except link-layer headers, in hexadecimal format.

Table 125: CLI monitor traffic Command Options (*continued*)

Option	Description
size bytes	(Optional) Displays the number of bytes for each packet that you specify. If a packet header exceeds this size, the displayed packet header is truncated. The default value is 96 .
brief	(Optional) Displays minimum packet header information. This is the default.
detail	(Optional) Displays packet header information in moderate detail. For some protocols, you must also use the size option to see detailed information.
extensive	(Optional) Displays the most extensive level of packet header information. For some protocols, you must also use the size option to see extensive information.

To quit the **monitor traffic** command and return to the command prompt, press Ctrl-C.

To limit the packet header information displayed by the **monitor traffic** command, include the **matching "expression"** option. An expression consists of one or more match conditions listed in [Table 126 on page 583](#), enclosed in quotation marks (" "). You can combine match conditions by using the logical operators listed in [Table 127 on page 585](#) (shown in order of highest to lowest precedence).

For example, to display TCP or UDP packet headers, enter:

```
user@host> monitor traffic matching "tcp || udp"
```

To compare the following types of expressions, use the relational operators listed in [Table 128 on page 585](#) (listed from highest to lowest precedence):

- Arithmetic—Expressions that use the arithmetic operators listed in [Table 128 on page 585](#).
- Binary—Expressions that use the binary operators listed in [Table 128 on page 585](#).
- Packet data accessor—Expressions that use the following syntax:

```
protocol [byte-offset <size>]
```

Replace **protocol** with any protocol in [Table 126 on page 583](#). Replace **byte-offset** with the byte offset, from the beginning of the packet header, to use for the comparison. The optional **size** parameter represents the number of bytes examined in the packet header—1, 2, or 4 bytes.

For example, the following command displays all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

Table 126: CLI monitor traffic Match Conditions

Match Condition	Description
Entity Type	

Table 126: CLI monitor traffic Match Conditions (*continued*)

Match Condition	Description
host [<i>address</i> <i>hostname</i>]	Matches packet headers that contain the specified address or hostname. You can prepend any of the following protocol match conditions, followed by a space, to host : arp , ip , rarp , or any of the Directional match conditions.
network address	Matches packet headers with source or destination addresses containing the specified network address.
network address mask <i>mask</i>	Matches packet headers containing the specified network address and subnet mask.
port [<i>port-number</i> <i>port-name</i>]	Matches packet headers containing the specified source or destination TCP or UDP port number or port name.
Directional	
destination	Matches packet headers containing the specified destination. Directional match conditions can be prepended to any Entity Type match conditions, followed by a space.
source	Matches packet headers containing the specified source.
source and destination	Matches packet headers containing the specified source <i>and</i> destination.
source or destination	Matches packet headers containing the specified source <i>or</i> destination.
Packet Length	
less bytes	Matches packets with lengths less than or equal to the specified value, in bytes.
greater bytes	Matches packets with lengths greater than or equal to the specified value, in bytes.
Protocol	
arp	Matches all ARP packets.
ether	Matches all Ethernet frames.
ether [broadcast multicast]	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with source or destination .
ether protocol [<i>address</i> (\arp \ip \rarp)	Matches Ethernet frames with the specified address or protocol type. The arguments arp , ip , and rarp are also independent match conditions, so they must be preceded with a backslash (\) when used in the ether protocol match condition.
icmp	Matches all ICMP packets.
ip	Matches all IP packets.
ip [broadcast multicast]	Matches broadcast or multicast IP packets.

Table 126: CLI monitor traffic Match Conditions (*continued*)

Match Condition	Description
ip protocol [<i>address</i> (\icmp igrp \tcp \udp)]	Matches IP packets with the specified address or protocol type. The arguments icmp , tcp , and udp are also independent match conditions, so they must be preceded with a backslash (\) when used in the ip protocol match condition.
isis	Matches all IS-IS routing messages.
rarp	Matches all RARP packets.
tcp	Matches all TCP packets.
udp	Matches all UDP packets.

Table 127: CLI monitor traffic Logical Operators

Logical Operator	Description
!	Logical NOT. If the first condition does not match, the next condition is evaluated.
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.
 	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
()	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

Table 128: CLI monitor traffic Arithmetic, Binary, and Relational Operators

Operator	Description
Arithmetic Operator	
+	Addition operator.
–	Subtraction operator.
/	Division operator.
Binary Operator	
&	Bitwise AND.
*	Bitwise exclusive OR.
 	Bitwise inclusive OR.

Table 128: CLI monitor traffic Arithmetic, Binary, and Relational Operators (*continued*)

Operator	Description
Relational Operator	
<=	A match occurs if the first expression is less than or equal to the second.
>=	A match occurs if the first expression is greater than or equal to the second.
<	A match occurs if the first expression is less than the second.
>	A match occurs if the first expression is greater than the second.
=	A match occurs if the first expression is equal to the second.
!=	A match occurs if the first expression is not equal to the second.

The following is sample output from the **monitor traffic** command:

```
user@host> monitor traffic count 4 matching "arp" detail
```

```
Listening on fe-0/0/0, capture size 96 bytes 15:04:16.276780 In arp who-has
193.1.1.1 tell host1.site2.net 15:04:16.376848 In arp who-has host2.site2.net
tell host1.site2.net 15:04:16.376887 In arp who-has 193.1.1.2 tell host1.site2.net
15:04:16.601923 In arp who-has 193.1.1.3 tell host1.site2.net
```

Related Documentation

- [Packet Capture Overview on page 565](#)
- [Using the J-Web Packet Capture Tool on page 586](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 580](#)
- [Example: Configuring Packet Capture on an Interface on page 572](#)

Using the J-Web Packet Capture Tool

Supported Platforms [SRX Series, vSRX](#)

You can use the J-Web packet capture diagnostic tool when you need to quickly capture and analyze router control traffic on a device. Packet capture on the J-Web user interface allows you to capture traffic destined for, or originating from, the Routing Engine. You can use the J-Web packet capture tool to compose expressions with various matching criteria to specify the packets that you want to capture. You can either choose to decode and view the captured packets in the J-Web user interface as they are captured, or save the captured packets to a file and analyze them offline using packet analyzers such as Ethereal. The J-Web packet capture tool does not capture transient traffic.

To capture transient traffic and entire IPv4 data packets for offline analysis, you must configure packet capture with the J-Web user interface or CLI configuration editor.

To use J-Web packet capture:

1. Select **Troubleshoot>Packet Capture**.
2. Enter information into the Packet Capture page (see [Table 129 on page 587](#)). The sample configuration captures the next 10 TCP packets originating from the IP address **10.1.40.48** on port 23 and passing through the Gigabit Ethernet interface **ge-0/0/0**.
3. Save the captured packets to a file, or specify other advanced options by clicking the expand icon next to Advanced options.
4. Click **Start**.

The captured packet headers are decoded and appear in the Packet Capture display.
5. Do one of the following:
 - To stop capturing the packets and stay on the same page while the decoded packet headers are being displayed, click **Stop Capturing**.
 - To stop capturing packets and return to the Packet Capture page, click **OK**.

Table 129: Packet Capture Field Summary

Field	Function	Your Action
Interface	Specifies the interface on which the packets are captured. If you select default , packets on the Ethernet management port 0 are captured.	Select an interface from the list—for example, ge-0/0/0 .
Detail level	Specifies the extent of details to be displayed for the packet headers. <ul style="list-style-type: none">• Brief—Displays the minimum packet header information. This is the default.• Detail—Displays packet header information in moderate detail.• Extensive—Displays the maximum packet header information.	Select Detail from the list.
Packets	Specifies the number of packets to be captured. Values range from 1 to 1000. Default is 10. Packet capture stops capturing packets after this number is reached.	Select the number of packets to be captured from the list—for example, 10.

Table 129: Packet Capture Field Summary (*continued*)

Field	Function	Your Action
Addresses	<p>Specifies the addresses to be matched for capturing the packets using a combination of the following parameters:</p> <ul style="list-style-type: none"> Direction—Matches the packet headers for IP address, hostname, or network address of the source, destination or both. Type—Specifies if packet headers are matched for host address or network address. <p>You can add multiple entries to refine the match criteria for addresses.</p>	<p>Select address-matching criteria. For example:</p> <ol style="list-style-type: none"> From the Direction list, select source. From the Type list, select host. In the Address box, type 10.1.40.48. Click Add.
Protocols	Matches the protocol for which packets are captured. You can choose to capture TCP, UDP, or ICMP packets or a combination of TCP, UDP, and ICMP packets.	Select a protocol from the list—for example, tcp .
Ports	Matches packet headers containing the specified source or destination TCP or UDP port number or port name.	<p>Select a direction and a port. For example:</p> <ol style="list-style-type: none"> From the Type list, select src. In the Port box, type 23.
Advanced Options		
Absolute TCP Sequence	Specifies that absolute TCP sequence numbers are to be displayed for the packet headers.	<ul style="list-style-type: none"> Display absolute TCP sequence numbers in the packet headers by selecting this check box. Stop displaying absolute TCP sequence numbers in the packet headers by clearing this check box.
Layer 2 Headers	Specifies that link-layer packet headers to display.	<ul style="list-style-type: none"> Include link-layer packet headers while capturing packets, by selecting this check box. Exclude link-layer packet headers while capturing packets by clearing this check box.
Non-Promiscuous	<p>Specifies not to place the interface in promiscuous mode, so that the interface reads only packets addressed to it.</p> <p>In promiscuous mode, the interface reads every packet that reaches it.</p>	<ul style="list-style-type: none"> Read all packets that reach the interface by selecting this check box. Read only packets addressed to the interface by clearing this check box.
Display Hex	Specifies that packet headers, except link-layer headers, are to be displayed in hexadecimal format.	<ul style="list-style-type: none"> Display the packet headers in hexadecimal format by selecting this check box. Stop displaying the packet headers in hexadecimal format by clearing this check box.
Display ASCII and Hex	Specifies that packet headers are to be displayed in hexadecimal and ASCII format.	<ul style="list-style-type: none"> Display the packet headers in ASCII and hexadecimal formats by selecting this check box. Stop displaying the packet headers in ASCII and hexadecimal formats by clearing this check box.

Table 129: Packet Capture Field Summary (*continued*)

Field	Function	Your Action
Header Expression	Specifies the match condition for the packets to capture. The match conditions you specify for Addresses, Protocols, and Ports appear in expression format in this field.	Enter match conditions in expression format or modify the expression composed from the match conditions you specified for Addresses, Protocols, and Ports. If you change the match conditions specified for Addresses, Protocols, and Ports again, packet capture overwrites your changes with the new match conditions.
Packet Size	Specifies the number of bytes to be displayed for each packet. If a packet header exceeds this size, the display is truncated for the packet header. The default value is 96 bytes.	Type the number of bytes you want to capture for each packet header—for example, 256 .
Don't Resolve Addresses	Specifies that IP addresses are not to be resolved into hostnames in the packet headers displayed.	<ul style="list-style-type: none"> Prevent packet capture from resolving IP addresses to hostnames by selecting this check box. Resolve IP addresses into hostnames by clearing this check box.
No Timestamp	Suppresses the display of packet header timestamps.	<ul style="list-style-type: none"> Stop displaying timestamps in the captured packet headers by selecting this check box. Display the timestamp in the captured packet headers by clearing this check box.
Write Packet Capture File	Writes the captured packets to a file in PCAP format in /var/tmp . The files are named with the prefix jweb-pcap and the extension .pcap . If you select this option, the decoded packet headers do not appear on the packet capture page.	<ul style="list-style-type: none"> Save the captured packet headers to a file by selecting this check box. Decode and display the packet headers on the J-Web page by clearing this check box.

- Related Documentation**
- [Packet Capture Overview on page 565](#)
 - [Diagnostic Tools Overview on page 8](#)
 - [J-Web Packet Capture Results and Output Summary on page 589](#)
 - [Using the J-Web Ping MPLS Tool on page 555](#)
 - [Using the J-Web Ping Host Tool on page 552](#)
 - [Using the J-Web Traceroute Tool on page 540](#)

J-Web Packet Capture Results and Output Summary

Supported Platforms [SRX Series, vSRX](#)

[Table 130 on page 590](#) summarizes the output in the packet capture display.

Table 130: J-Web Packet Capture Results and Output Summary

Field	Description
<i>timestamp</i>	<p>Time when the packet was captured. The timestamp 00:45:40.823971 means 00 hours (12.00 a.m.), 45 minutes, and 40.823971 seconds.</p> <p>NOTE: The time displayed is local time.</p>
<i>direction</i>	<p>Direction of the packet. Specifies whether the packet originated from the Routing Engine (Out), or was destined for the Routing Engine (In).</p>
<i>protocol</i>	<p>Protocol for the packet.</p> <p>In the sample output, IP indicates the Layer 3 protocol.</p>
<i>source address</i>	<p>Hostname, if available, or IP address and the port number of the packet's origin. If the Don't Resolve Addresses check box is selected, only the IP address of the source displays.</p> <p>NOTE: When a string is defined for the port, the packet capture output displays the string instead of the port number.</p>
<i>destination address</i>	<p>Hostname, if available, or IP address of the packet's destination with the port number. If the Don't Resolve Addresses check box is selected, only the IP address of the destination and the port appear.</p> <p>NOTE: When a string is defined for the port, the packet capture output displays the string instead of the port number.</p>
<i>protocol</i>	<p>Protocol for the packet.</p> <p>In the sample output, TCP indicates the Layer 4 protocol.</p>
<i>data size</i>	<p>Size of the packet (in bytes).</p>

**Related
Documentation**

- [Packet Capture Overview on page 565](#)
- [Diagnostic Tools Overview on page 8](#)
- [Using the J-Web Packet Capture Tool on page 586](#)

Troubleshooting Security Devices

- Recovering the Root Password for SRX Series Devices on page 591
- Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only) on page 593
- Troubleshooting the Link Services Interface on page 593
- Troubleshooting Security Policies on page 604
- Understanding Log Error Messages for Troubleshooting ISSU-Related Problems on page 606

Recovering the Root Password for SRX Series Devices

Supported Platforms [SRX Series, vSRX](#)

If you forget the root password for an SRX Series device, you can use the password recovery procedure to reset the root password. This procedure also involves disabling the watchdog functionality to allow the system to properly boot into single-user mode (KB article 17565).



NOTE: You need console access to recover the root password

To recover the root password for an SRX Series device:

1. Power on the device by pressing the power button on the front panel. Verify that the **POWER** LED on the front panel turns green.

The device's boot sequence on your management device appears on the terminal emulation screen.
2. When the autoboot completes, press the Spacebar a few times to access the bootstrap loader prompt.
3. In operational mode, disable the watchdog functionality and enter **boot -s** to start up the system in single-user mode.

```
loader>boot -s
```



.....

NOTE: For SRX1500 Series device, the following prompt appears instead of loader prompt.

OK boot -s.

.....

The SRX Series device will start up in single-user mode.

4. Enter **recovery** to start the root password recovery procedure.

System watchdog timer disabled.

Enter full pathname of shell or 'recovery' for root password recovery or RETURN for /bin/sh: **recovery**

5. Enter configuration mode in the CLI.

6. Set the root password.

[edit]

user@host# **set system root-authentication plain-text-password**

7. Enter the new root password.

New password: **juniper1**

Retype new password:

8. At the second prompt, reenter the new root password.

9. If you are finished configuring the network, commit the configuration.

root@host# **commit**

commit complete

10. Exit from configuration mode.

11. Exit from operational mode.

12. Enter **y** to reboot the device.

Reboot the system? [y/n] **y**

The start up messages display on the screen.

13. Once again, press the Spacebar a few times to access the bootstrap loader prompt.

14. In operational mode, enable the watchdog functionality and enter **boot** to start up the system.

loader>**watchdog enable**

loader>**boot**

15. The SRX Series device starts up again and prompts you to enter a user name and password. Enter the newly configured password:

```
Wed Jul 12 14:20:21 UTC 2011
Deviceabc (ttyu0)
login: root
Password: juniper1
```

Related Documentation • [System Log Messages](#)

Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only)

Supported Platforms [SRX5400](#), [SRX5600](#), [SRX5800](#), [vSRX](#)

Problem **Description:** The address of a hostname in an address book entry that is used in a security policy might fail to resolve correctly.

Cause Normally, address book entries that contain dynamic hostnames refresh automatically for SRX Series devices. The TTL field associated with a DNS entry indicates the time after which the entry should be refreshed in the policy cache. Once the TTL value expires, the SRX Series device automatically refreshes the DNS entry for an address book entry.

However, if the SRX Series device is unable to obtain a response from the DNS server (for example, the DNS request or response packet is lost in the network or the DNS server cannot send a response), the address of a hostname in an address book entry might fail to resolve correctly. This can cause traffic to drop as no security policy or session match is found.

Solution The master administrator can use the **show security dns-cache** command to display DNS cache information on the SRX Series device. If the DNS cache information needs to be refreshed, the master administrator can use the **clear security dns-cache** command.



NOTE: These commands are only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Related Documentation • [Understanding Logical System Security Policies](#)

Troubleshooting the Link Services Interface

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [vSRX](#)

To solve configuration problems on a link services interface:

- [Determine Which CoS Components Are Applied to the Constituent Links on page 594](#)
- [Determine What Causes Jitter and Latency on the Multilink Bundle on page 596](#)
- [Determine If LFI and Load Balancing Are Working Correctly on page 596](#)
- [Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device on page 603](#)

Determine Which CoS Components Are Applied to the Constituent Links

Problem **Description:** You are configuring a multilink bundle, but you also have traffic without MLPPP encapsulation passing through constituent links of the multilink bundle. Do you apply all CoS components to the constituent links, or is applying them to the multilink bundle enough?

Solution You can apply a scheduler map to the multilink bundle and its constituent links. Although you can apply several CoS components with the scheduler map, configure only the ones that are required. We recommend that you keep the configuration on the constituent links simple to avoid unnecessary delay in transmission.

[Table 131 on page 594](#) shows the CoS components to be applied on a multilink bundle and its constituent links.

Table 131: CoS Components Applied on Multilink Bundles and Constituent Links

Cos Component	Multilink Bundle	Constituent Links	Explanation
Classifier	Yes	No	CoS classification takes place on the incoming side of the interface, not on the transmitting side, so no classifiers are needed on constituent links.
Forwarding class	Yes	No	Forwarding class is associated with a queue, and the queue is applied to the interface by a scheduler map. The queue assignment is predetermined on the constituent links. All packets from Q2 of the multilink bundle are assigned to Q2 of the constituent link, and packets from all the other queues are queued to Q0 of the constituent link.

Table 131: CoS Components Applied on Multilink Bundles and Constituent Links (*continued*)

Cos Component	Multilink Bundle	Constituent Links	Explanation
Scheduler map	Yes	Yes	<p>Apply scheduler maps on the multilink bundle and the constituent link as follows:</p> <ul style="list-style-type: none"> • Transmit rate—Make sure that the relative order of the transmit rate configured on Q0 and Q2 is the same on the constituent links as on the multilink bundle. • Scheduler priority—Make sure that the relative order of the scheduler priority configured on Q0 and Q2 is the same on the constituent links as on the multilink bundle. • Buffer size—Because all non-LFI packets from the multilink bundle transit on Q0 of the constituent links, make sure that the buffer size on Q0 of the constituent links is large enough. • RED drop profile—Configure a RED drop profile on the multilink bundle only. Configuring the RED drop profile on the constituent links applies a back pressure mechanism that changes the buffer size and introduces variation. Because this behavior might cause fragment drops on the constituent links, make sure to leave the RED drop profile at the default settings on the constituent links.
Shaping rate for a per-unit scheduler or an interface-level scheduler	No	Yes	Because per-unit scheduling is applied only at the end point, apply this shaping rate to the constituent links only. Any configuration applied earlier is overwritten by the constituent link configuration.
Transmit-rate exact or queue-level shaping	Yes	No	The interface-level shaping applied on the constituent links overrides any shaping on the queue. Thus apply transmit-rate exact shaping on the multilink bundle only.
Rewrite rules	Yes	No	Rewrite bits are copied from the packet into the fragments automatically during fragmentation. Thus what you configure on the multilink bundle is carried on the fragments to the constituent links.
Virtual channel group	Yes	No	Virtual channel groups are identified through firewall filter rules that are applied on packets only before the multilink bundle. Thus you do not need to apply the virtual channel group configuration to the constituent links.

- See Also**
- [Link Services Interfaces Overview](#)
 - [Troubleshooting the Link Services Interface on page 593](#)
 - [See the Junos OS Class of Service Configuration Guide for Security Devices](#)

Determine What Causes Jitter and Latency on the Multilink Bundle

Problem **Description:** To test jitter and latency, you send three streams of IP packets. All packets have the same IP precedence settings. After configuring LFI and CRTP, the latency increased even over a noncongested link. How can you reduce jitter and latency?

Solution To reduce jitter and latency, do the following:

1. Make sure that you have configured a shaping rate on each constituent link.
2. Make sure that you have not configured a shaping rate on the link services interface.
3. Make sure that the configured shaping rate value is equal to the physical interface bandwidth.
4. If shaping rates are configured correctly, and jitter still persists, contact the Juniper Networks Technical Assistance Center (JTAC).

See Also • [RPM Overview on page 341](#)

Determine If LFI and Load Balancing Are Working Correctly

Problem **Description:** In this case, you have a single network that supports multiple services. The network transmits data and delay-sensitive voice traffic. After configuring MLPPP and LFI, make sure that voice packets are transmitted across the network with very little delay and jitter. How can you find out if voice packets are being treated as LFI packets and load balancing is performed correctly?

Solution When LFI is enabled, data (non-LFI) packets are encapsulated with an MLPPP header and fragmented to packets of a specified size. The delay-sensitive, voice (LFI) packets are PPP-encapsulated and interleaved between data packet fragments. Queuing and load balancing are performed differently for LFI and non-LFI packets.

To verify that LFI is performed correctly, determine that packets are fragmented and encapsulated as configured. After you know whether a packet is treated as an LFI packet or a non-LFI packet, you can confirm whether the load balancing is performed correctly.

Solution Scenario—Suppose two Juniper Networks devices, R0 and R1, are connected by a multilink bundle `lsq-0/0/0.0` that aggregates two serial links, `se-1/0/0` and `se-1/0/1`. On R0 and R1, MLPPP and LFI are enabled on the link services interface and the fragmentation threshold is set to 128 bytes.

In this example, we used a packet generator to generate voice and data streams. You can use the packet capture feature to capture and analyze the packets on the incoming interface.

The following two data streams were sent on the multilink bundle:

- 100 data packets of 200 bytes (larger than the fragmentation threshold)
- 500 data packets of 60 bytes (smaller than the fragmentation threshold)

The following two voice streams were sent on the multilink bundle:

- 100 voice packets of 200 bytes from source port 100
- 300 voice packets of 200 bytes from source port 200

To confirm that LFI and load balancing are performed correctly:



NOTE: Only the significant portions of command output are displayed and described in this example.

1. Verify packet fragmentation. From operational mode, enter the **show interfaces lsq-0/0/0** command to check that large packets are fragmented correctly.

```
user@R0#> show interfaces lsq-0/0/0
Physical interface: lsq-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 29
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2006-08-01 10:45:13 PDT (2w0d 06:06 ago)
  Input rate      : 0 bps (0 pps)
  Output rate     : 0 bps (0 pps)

Logical interface lsq-0/0/0.0 (Index 69) (SNMP ifIndex 42)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  Bandwidth: 16mbps
  Statistics
  Bundle:
    Fragments:
      Input :      0      0      0      0
      Output:    1100      0    118800    0
    Packets:
      Input :      0      0      0      0
      Output:    1000      0    112000    0
  ...
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 9.9.9/24, Local: 9.9.9.10
```

Meaning—The output shows a summary of packets transiting the device on the multilink bundle. Verify the following information on the multilink bundle:

- The total number of transiting packets = 1000
- The total number of transiting fragments=1100
- The number of data packets that were fragmented =100

The total number of packets sent (600 + 400) on the multilink bundle match the number of transiting packets (1000), indicating that no packets were dropped.

The number of transiting fragments exceeds the number of transiting packets by 100, indicating that 100 large data packets were correctly fragmented.

Corrective Action—If the packets are not fragmented correctly, check your fragmentation threshold configuration. Packets smaller than the specified fragmentation threshold are not fragmented.

2. Verify packet encapsulation. To find out whether a packet is treated as an LFI or

non-LFI packet, determine its encapsulation type. LFI packets are PPP encapsulated, and non-LFI packets are encapsulated with both PPP and MLPPP. PPP and MLPPP encapsulations have different overheads resulting in different-sized packets. You can compare packet sizes to determine the encapsulation type.

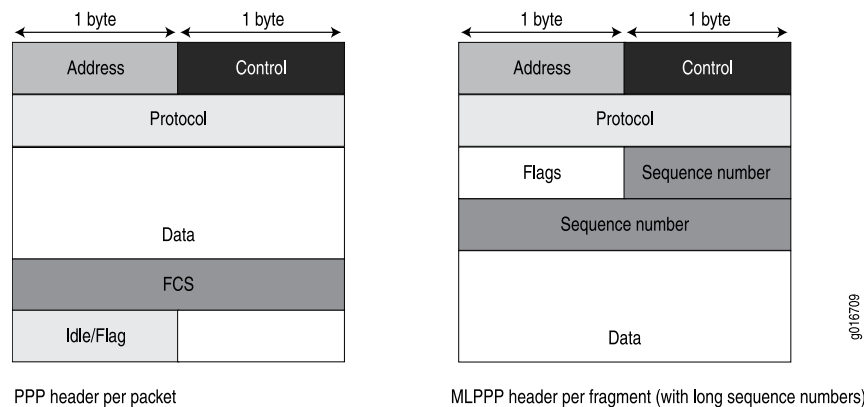
A small unfragmented data packet contains a PPP header and a single MLPPP header. In a large fragmented data packet, the first fragment contains a PPP header and an MLPPP header, but the consecutive fragments contain only an MLPPP header.

PPP and MLPPP encapsulations add the following number of bytes to a packet:

- PPP encapsulation adds 7 bytes:
4 bytes of header+2 bytes of frame check sequence (FCS)+1 byte that is idle or contains a flag
- MLPPP encapsulation adds between 6 and 8 bytes:
4 bytes of PPP header+2 to 4 bytes of multilink header

Figure 12 on page 599 shows the overhead added to PPP and MLPPP headers.

Figure 12: PPP and MLPPP Headers



For CRTP packets, the encapsulation overhead and packet size are even smaller than for an LFI packet. For more information, see [Example: Configuring the Compressed Real-Time Transport Protocol](#).

Table 132 on page 599 shows the encapsulation overhead for a data packet and a voice packet of 70 bytes each. After encapsulation, the size of the data packet is larger than the size of the voice packet.

Table 132: PPP and MLPPP Encapsulation Overhead

Packet Type	Encapsulation	Initial Packet Size	Encapsulation Overhead	Packet Size after Encapsulation
Voice packet (LFI)	PPP	70 bytes	4 + 2 + 1 = 7 bytes	77 bytes
Data fragment (non-LFI) with short sequence	MLPPP	70 bytes	4 + 2 + 1 + 4 + 2 = 13 bytes	83 bytes

Table 132: PPP and MLPPP Encapsulation Overhead (*continued*)

Packet Type	Encapsulation	Initial Packet Size	Encapsulation Overhead	Packet Size after Encapsulation
Data fragment (non-LFI) with long sequence	MLPPP	70 bytes	4 + 2 + 1 + 4 + 4 = 15 bytes	85 bytes

From operational mode, enter the **show interfaces queue** command to display the size of transmitted packet on each queue. Divide the number of bytes transmitted by the number of packets to obtain the size of the packets and determine the encapsulation type.

3. Verify load balancing. From operational mode, enter the **show interfaces queue** command on the multilink bundle and its constituent links to confirm whether load balancing is performed accordingly on the packets.

```

user@R0> show interfaces queue lsq-0/0/0
Physical interface: lsq-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 29
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: DATA
  Queued:
    Packets      :           600      0 pps
    Bytes        :          44800      0 bps
  Transmitted:
    Packets      :           600      0 pps
    Bytes        :          44800      0 bps
    Tail-dropped packets :           0      0 pps
    RED-dropped packets  :           0      0 pps
  ...
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0      0 pps
    Bytes        :              0      0 bps
  ...
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets      :           400      0 pps
    Bytes        :          61344      0 bps
  Transmitted:
    Packets      :           400      0 pps
    Bytes        :          61344      0 bps
  ...
Queue: 3, Forwarding classes: NC
  Queued:
    Packets      :              0      0 pps
    Bytes        :              0      0 bps
  ...

user@R0> show interfaces queue se-1/0/0
Physical interface: se-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 35
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use

```

```

Queue: 0, Forwarding classes: DATA
  Queued:
    Packets      :           350      0 pps
    Bytes        :        24350      0 bps
  Transmitted:
    Packets      :           350      0 pps
    Bytes        :        24350      0 bps
  ...
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0      0 pps
    Bytes        :              0      0 bps
  ...
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets      :           100      0 pps
    Bytes        :        15272      0 bps
  Transmitted:
    Packets      :           100      0 pps
    Bytes        :        15272      0 bps
  ...
Queue: 3, Forwarding classes: NC
  Queued:
    Packets      :             19      0 pps
    Bytes        :           247      0 bps
  Transmitted:
    Packets      :             19      0 pps
    Bytes        :           247      0 bps
  ...

```

user@R0> show interfaces queue se-1/0/1

Physical interface: se-1/0/1, Enabled, Physical link is Up

Interface index: 142, SNMP ifIndex: 38

Forwarding classes: 8 supported, 8 in use

Egress queues: 8 supported, 8 in use

```

Queue: 0, Forwarding classes: DATA
  Queued:
    Packets      :           350      0 pps
    Bytes        :        24350      0 bps
  Transmitted:
    Packets      :           350      0 pps
    Bytes        :        24350      0 bps
  ...
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0      0 pps
    Bytes        :              0      0 bps
  ...
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets      :           300      0 pps
    Bytes        :        45672      0 bps
  Transmitted:
    Packets      :           300      0 pps
    Bytes        :        45672      0 bps
  ...
Queue: 3, Forwarding classes: NC
  Queued:
    Packets      :             18      0 pps
    Bytes        :           234      0 bps
  Transmitted:

```

```

Packets          :          18          0 pps
Bytes            :          234          0 bps

```

Meaning—The output from these commands shows the packets transmitted and queued on each queue of the link services interface and its constituent links.

[Table 133 on page 602](#) shows a summary of these values. (Because the number of transmitted packets equaled the number of queued packets on all the links, this table shows only the queued packets.)

Table 133: Number of Packets Transmitted on a Queue

Packets Queued	Bundle lsq-0/0/0.0	Constituent Link se-1/0/0	Constituent Link se-1/0/1	Explanation
Packets on Q0	600	350	350	The total number of packets transiting the constituent links (350+350 = 700) exceeded the number of packets queued (600) on the multilink bundle.
Packets on Q2	400	100	300	The total number of packets transiting the constituent links equaled the number of packets on the bundle.
Packets on Q3	0	19	18	The packets transiting Q3 of the constituent links are for keepalive messages exchanged between constituent links. Thus no packets were counted on Q3 of the bundle.

On the multilink bundle, verify the following:

- The number of packets queued matches the number transmitted. If the numbers match, no packets were dropped. If more packets were queued than were transmitted, packets were dropped because the buffer was too small. The buffer size on the constituent links controls congestion at the output stage. To correct this problem, increase the buffer size on the constituent links.
- The number of packets transiting Q0 (600) matches the number of large and small data packets received (100+500) on the multilink bundle. If the numbers match, all data packets correctly transited Q0.
- The number of packets transiting Q2 on the multilink bundle (400) matches the number of voice packets received on the multilink bundle. If the numbers match, all voice LFI packets correctly transited Q2.

On the constituent links, verify the following:

- The total number of packets transiting Q0 (350+350) matches the number of data packets and data fragments (500+200). If the numbers match, all the data packets after fragmentation correctly transited Q0 of the constituent links.

Packets transited both constituent links, indicating that load balancing was correctly performed on non-LFI packets.

- The total number of packets transiting Q2 (300+100) on constituent links matches the number of voice packets received (400) on the multilink bundle. If the numbers match, all voice LFI packets correctly transited Q2.

LFI packets from source port **100** transited **se-1/0/0**, and LFI packets from source port **200** transited **se-1/0/1**. Thus all LFI (Q2) packets were hashed based on the source port and correctly transited both constituent links.

Corrective Action—If the packets transited only one link, take the following steps to resolve the problem:

- Determine whether the physical link is **up** (operational) or **down** (unavailable). An unavailable link indicates a problem with the PIM, interface port, or physical connection (link-layer errors). If the link is operational, move to the next step.
 - Verify that the classifiers are correctly defined for non-LFI packets. Make sure that non-LFI packets are not configured to be queued to Q2. All packets queued to Q2 are treated as LFI packets.
 - Verify that at least one of the following values is different in the LFI packets: source address, destination address, IP protocol, source port, or destination port. If the same values are configured for all LFI packets, the packets are all hashed to the same flow and transit the same link.
- Use the results to verify load balancing.

- See Also**
- [Link Services Interfaces Overview](#)
 - [Troubleshooting the Link Services Interface on page 593](#)

Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device

Problem **Description:** You are configuring a permanent virtual circuit (PVC) between T1, E1, T3, or E3 interfaces on a Juniper Networks device and a third-party device, and packets are being dropped and ping fails.

Solution If the third-party device does not have the same FRF.12 support as the Juniper Networks device or supports FRF.12 in a different way, the Juniper Networks device interface on the PVC might discard a fragmented packet containing FRF.12 headers and count it as a "Policed Discard."

As a workaround, configure multilink bundles on both peers, and configure fragmentation thresholds on the multilink bundles.

- See Also**
- [Link Services Interfaces Overview](#)

- [Troubleshooting the Link Services Interface on page 593](#)

Troubleshooting Security Policies

Supported Platforms [SRX Series, vSRX](#)

- [Synchronizing Policies Between Routing Engine and Packet Forwarding Engine on page 604](#)
- [Checking a Security Policy Commit Failure on page 604](#)
- [Verifying a Security Policy Commit on page 605](#)
- [Debugging Policy Lookup on page 606](#)

Synchronizing Policies Between Routing Engine and Packet Forwarding Engine

Problem **Description:** Security policies are stored in both the Routing Engine and the Packet Forwarding Engine. After you modify a policy, you commit the configuration on the Routing Engine, and it is synchronized to the Packet Forwarding Engine.

Environment: The policies in the Routing Engine and Packet Forwarding Engine must be in sync for the configuration to be committed. However, under certain circumstances, policies in the Routing Engine and the Packet Forwarding Engine might be out of sync, which causes the commit to fail.

Symptoms: The following error message appears if you attempt to commit a configuration when the policies in the Routing Engine and Packet Forwarding Engine are out of sync:
Policy is out of sync between RE and PFE <SPU-name(s)> Please resync before commit.

Solution Synchronize the policies as follows:

- Reboot the device (standalone)
- Reboot the devices (chassis cluster)

See Also

- [Verifying a Security Policy Commit on page 605](#)
- [Checking a Security Policy Commit Failure on page 604](#)
- [Debugging Policy Lookup on page 606](#)
- [Monitoring Policy Statistics on page 453](#)

Checking a Security Policy Commit Failure

Problem **Description:** Most policy configuration failures occur during a commit or runtime. Commit failures are reported directly on the CLI when you execute the CLI command **commit-check** in configuration mode. These errors are configuration errors, and you cannot commit the configuration without fixing these errors.

Solution To fix these errors, do the following:

1. Review your configuration data.
2. Open the file `/var/log/nsd_chk_only`. This file is overwritten each time you perform a commit check and contains detailed failure information.

See Also

- [Synchronizing a Security Policy on SRX Series Devices](#)
- [Verifying a Security Policy Commit on page 605](#)
- [Debugging Policy Lookup on page 606](#)
- [Monitoring Policy Statistics on page 453](#)

Verifying a Security Policy Commit

Problem Description: Upon performing a policy configuration commit, if you notice that the system behavior is incorrect, use the following steps to troubleshoot this problem:

Solution

1. **Operational `show` Commands**—Execute the operational commands for security policies and verify that the information shown in the output is consistent with what you expected. If not, the configuration needs to be changed appropriately.
2. **Traceoptions**—Set the **traceoptions** command in your policy configuration. The flags under this hierarchy can be selected as per user analysis of the **show** command output. If you cannot determine what flag to use, the flag option **all** can be used to capture all trace logs.

```
user@host# set security policies traceoptions <flag all>
```

You can also configure an optional filename to capture the logs.

```
user@host# set security policies traceoptions <filename>
```

If you specified a filename in the trace options, you can look in the `/var/log/<filename>` for the log file to ascertain if any errors were reported in the file. (If you did not specify a filename, the default filename is `eventd`.) The error messages indicate the place of failure and the appropriate reason.

After configuring the trace options, you must recommit the configuration change that caused the incorrect system behavior.

See Also

- [Synchronizing a Security Policy on SRX Series Devices](#)
- [Checking a Security Policy Commit Failure on page 604](#)
- [Debugging Policy Lookup on page 606](#)
- [Monitoring Policy Statistics on page 453](#)

Debugging Policy Lookup

Problem **Description:** When you have the correct configuration, but some traffic was incorrectly dropped or permitted, you can enable the **lookup** flag in the security policies traceoptions. The **lookup** flag logs the lookup related traces in the trace file.

Solution `user@host# set security policies traceoptions <flag lookup>`

See Also

- [Checking a Security Policy Commit Failure on page 604](#)
- [Verifying a Security Policy Commit on page 605](#)
- [Monitoring Policy Statistics on page 453](#)

Related Documentation

- [Synchronizing Policies Between Routing Engine and Packet Forwarding Engine on page 604](#)
- [Security Policies Overview](#)
- [Checking a Security Policy Commit Failure on page 604](#)
- [Verifying a Security Policy Commit on page 605](#)
- [Debugging Policy Lookup on page 606](#)
- [Monitoring Policy Statistics on page 453](#)

Understanding Log Error Messages for Troubleshooting ISSU-Related Problems

Supported Platforms SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, vSRX

The following problems might occur during an ISSU upgrade. You can identify the errors by using the details in the logs. You can also see the details of the error messages in the System Log Explorer.

- [Chassisd Process Errors on page 606](#)
- [Kernel State Synchronization on page 607](#)
- [Installation-Related Errors on page 607](#)
- [ISSU Support-Related Errors on page 607](#)
- [Redundancy Group Failover Errors on page 608](#)
- [Initial Validation Checks Fail on page 608](#)
- [Understanding Common Error Handling for ISSU on page 609](#)

Chassisd Process Errors

Problem **Description:** Errors related to chassisd.

Solution Use the error messages to understand the issues related to chassisd.

When ISSU starts, a request is sent to chassisd to check whether there are any problems related to the ISSU from a chassis perspective. If there is a problem, a log message is created.

Kernel State Synchronization

Problem **Description:** Errors related to ksyncd.

Solution Use the following error messages to understand the issues related to ksyncd:

```
Failed to get kernel-replication error information from Standby Routing Engine.  
mgd_slave_peer_has_errors() returns error at line 4414 in mgd_package_issu.
```

ISSU checks whether there are any ksyncd errors on the secondary node (node 1) and displays the error message if there are any problems and aborts the upgrade.

Installation-Related Errors

Problem **Description:** The install image file does not exist or the remote site is inaccessible.

Solution Use the following error messages to understand the installation-related problems:

```
error: File does not exist: /var/tmp/junos-srx5000-11.4X3.2-domest  
error: Couldn't retrieve package /var/tmp/junos-srx5000-11.4X3.2-domest
```

ISSU downloads the install image as specified in the ISSU command as an argument. The image file can be a local file or located at a remote site. If the file does not exist or the remote site is inaccessible, an error is reported.

ISSU Support-Related Errors

Problem **Description:** Installation failure occurs because of unsupported software and unsupported feature configuration.

Solution Use the following error messages to understand the compatibility-related problems:

```
WARNING: Current configuration not compatible with  
/var/tmp/junos-srx5000-11.4X3.2-domestic.tgz  
Exiting in-service-upgrade window  
Exiting in-service-upgrade window
```

Redundancy Group Failover Errors

Problem **Description:** Problem with automatic redundancy group (RG) failure.

Solution Use the following error messages to understand the problem:

```
failover all RG 1+ groups to node 0
error: Command failed. None of the redundancy-groups has been failed over.
      Some redundancy-groups on node1 are already in manual failover mode.
      Please execute 'failover reset all' first..
```

Initial Validation Checks Fail

Problem **Description:** The initial validation checks fail.

Solution The validation checks fail if the image is not present or if the image file is corrupt. The following error messages are displayed when initial validation checks fail when the image is not present and the ISSU is aborted:

When Image Is Not Present

```
user@host> ...0120914_srx_12q1_major2.2-539764-domestic.tgz reboot
Chassis ISSU Started
Chassis ISSU Started
ISSU: Validating Image
Initiating in-service-upgrade
Initiating in-service-upgrade
Fetching package...
error: File does not exist:
/var/tmp/junos-srx1k3k-12.1I20120914_srx_12q1_major2.2-539764-domestic.tgz
error: Couldn't retrieve package
/var/tmp/junos-srx1k3k-12.1I20120914_srx_12q1_major2.2-539764-domestic.tgz
Exiting in-service-upgrade window
Exiting in-service-upgrade window
Chassis ISSU Aborted
Chassis ISSU Aborted
Chassis ISSU Aborted
ISSU: IDLE
ISSU aborted; exiting ISSU window.
```

When Image File Is Corrupted

If the image file is corrupted, the following output displays:

```
user@host> ...junos-srx1k3k-11.4X9-domestic.tgz_1 reboot
Chassis ISSU Started
node1:
-----
Chassis ISSU Started
ISSU: Validating Image
Initiating in-service-upgrade
```

```

node1:
-----
Initiating in-service-upgrade
ERROR: Cannot use /var/tmp/junos-srx1k3k-11.4X9-domestic.tgz_1:
gzip: stdin: invalid compressed data--format violated
tar: Child returned status 1
tar: Error exit delayed from previous errors
ERROR: It may have been corrupted during download.
ERROR: Please try again, making sure to use a binary transfer.
Exiting in-service-upgrade window

node1:
-----
Exiting in-service-upgrade window
Chassis ISSU Aborted
Chassis ISSU Aborted

node1:
-----
Chassis ISSU Aborted
ISSU: IDLE
ISSU aborted; exiting ISSU window.

{primary:node0}

```

The primary node validates the device configuration to ensure that it can be committed using the new software version. If anything goes wrong, the ISSU aborts and error messages are displayed.

Understanding Common Error Handling for ISSU

Problem **Description:** You might encounter some problems in the course of an ISSU. This section provides details on how to handle them.

Solution Any errors encountered during an ISSU result in the creation of log messages, and ISSU continues to function without impact to traffic. If reverting to previous versions is required, the event is either logged or the ISSU is halted, so as not to create any mismatched versions on both nodes of the chassis cluster. [Table 134 on page 609](#) provides some of the common error conditions and the workarounds for them. The sample messages used in the [Table 134 on page 609](#) are from the SRX1500 device and are also applicable to all supported SRX Series devices.

Table 134: ISSU-Related Errors and Solutions

Error Conditions	Solutions
Attempt to initiate an ISSU when previous instance of an ISSU is already in progress	<p>The following message is displayed:</p> <p>warning: ISSU in progress</p> <p>You can abort the current ISSU process, and initiate the ISSU again using the request chassis cluster in-service-upgrade abort command.</p>

Table 134: ISSU-Related Errors and Solutions (*continued*)

Error Conditions	Solutions
Reboot failure on the secondary node	<p>No service downtime occurs, because the primary node continues to provide required services. Detailed console messages are displayed requesting that you manually clear existing ISSU states and restore the chassis cluster.</p> <pre>error: [Oct 6 12:30:16]: Reboot secondary node failed (error-code: 4.1)</pre> <pre>error: [Oct 6 12:30:16]: ISSU Aborted! Backup node maybe in inconsistent state, Please restore backup node [Oct 6 12:30:16]: ISSU aborted. But, both nodes are in ISSU window. Please do the following: 1. Rollback the node with the newer image using rollback command Note: use the 'node' option in the rollback command otherwise, images on both nodes will be rolled back 2. Make sure that both nodes (will) have the same image 3. Ensure the node with older image is primary for all RGs 4. Abort ISSU on both nodes 5. Reboot the rolled back node</pre>
Secondary node failed to complete the cold synchronization	<p>The primary node times out if the secondary node fails to complete the cold synchronization. Detailed console messages are displayed that you manually clear existing ISSU states and restore the chassis cluster. No service downtime occurs in this scenario.</p> <pre>[Oct 3 14:00:46]: timeout waiting for secondary node node1 to sync(error-code: 6.1) Chassis control process started, pid 36707</pre> <pre>error: [Oct 3 14:00:46]: ISSU Aborted! Backup node has been upgraded, Please restore backup node [Oct 3 14:00:46]: ISSU aborted. But, both nodes are in ISSU window. Please do the following: 1. Rollback the node with the newer image using rollback command Note: use the 'node' option in the rollback command otherwise, images on both nodes will be rolled back 2. Make sure that both nodes (will) have the same image 3. Ensure the node with older image is primary for all RGs 4. Abort ISSU on both nodes 5. Reboot the rolled back node</pre>

Table 134: ISSU-Related Errors and Solutions (*continued*)

Error Conditions	Solutions
Failover of newly upgraded secondary failed	<p>No service downtime occurs, because the primary node continues to provide required services. Detailed console messages are displayed requesting that you manually clear existing ISSU states and restore the chassis cluster.</p> <pre>[Aug 27 15:28:17]: Secondary node0 ready for failover. [Aug 27 15:28:17]: Failing over all redundancy-groups to node0 ISSU: Preparing for Switchover error: remote rg1 priority zero, abort failover. [Aug 27 15:28:17]: failover all RGs to node node0 failed (error-code: 7.1) error: [Aug 27 15:28:17]: ISSU Aborted! [Aug 27 15:28:17]: ISSU aborted. But, both nodes are in ISSU window. Please do the following: 1. Rollback the node with the newer image using rollback command Note: use the 'node' option in the rollback command otherwise, images on both nodes will be rolled back 2. Make sure that both nodes (will) have the same image 3. Ensure the node with older image is primary for all RGs 4. Abort ISSU on both nodes 5. Reboot the rolled back node {primary:node1}</pre>
Upgrade failure on primary	No service downtime occurs, because the secondary node fails over as primary and continues to provide required services.
Reboot failure on primary node	<p>Before the reboot of the primary node, devices being out of the ISSU setup, no ISSU-related error messages are displayed. The following reboot error message is displayed if any other failure is detected:</p> <pre>Reboot failure on Before the reboot of primary node, devices will be out of ISSU setup and no primary node error messages will be displayed. Primary node</pre>

**Related
Documentation**

- *Troubleshooting Chassis Cluster ISSU-Related Problems*
- *Upgrading Both Devices in a Chassis Cluster Using an ISSU*
- *ISSU System Requirements*
- *Understanding the ISSU Process on Devices in a Chassis Cluster*

PART 10

Configuration Statements and Operational Commands

- [Configuration Statements: Accounting Options, Source Class Usage and Destination Class Usage Options on page 615](#)
- [Configuration Statements: Chassis Cluster on page 663](#)
- [Configuration Statements: Datapath Debug on page 671](#)
- [Configuration Statements: Health Monitoring on page 683](#)
- [Configuration Statements: Remote Monitoring \(RMON\) on page 691](#)
- [Configuration Statements: Resource Monitoring for Memory Regions on page 709](#)
- [Configuration Statements: Security Alarms on page 723](#)
- [Configuration Statements: SNMP on page 725](#)
- [Configuration Statements: SNMPv3 on page 765](#)
- [Configuration Statements: TWAMP on page 821](#)
- [Operational Commands on page 827](#)

Configuration Statements: Accounting Options, Source Class Usage and Destination Class Usage Options

- [accounting-options](#) on page 617
- [allow-clear](#) (Accounting Options) on page 619
- [archive-sites](#) on page 620
- [backup-on-failure](#) (Accounting Options) on page 621
- [class-usage-profile](#) on page 622
- [cleanup-interval](#) (Accounting Options) on page 623
- [compress](#) (Accounting Options) on page 624
- [counters](#) on page 624
- [destination-classes](#) on page 625
- [egress-stats](#) (Flat-File Accounting Options) on page 626
- [fields](#) (Flat-File Accounting Options) on page 628
- [fields](#) (for Interface Profiles) on page 630
- [fields](#) (for Routing Engine Profiles) on page 631
- [file](#) (Associating with a Profile) on page 632
- [file](#) (Configuring a Log File) on page 633
- [file](#) (Flat-File Accounting Options) on page 634
- [files](#) on page 635
- [filter-profile](#) on page 636
- [flat-file-profile](#) (Accounting Options) on page 637
- [format](#) (Flat-File Accounting Options) on page 639
- [general-param](#) (Flat-File Accounting Options) on page 640
- [ingress-stats](#) (Flat-File Accounting Options) on page 642
- [interface-profile](#) on page 643
- [interval](#) on page 644
- [interval](#) (Flat-File Accounting Options) on page 645

- [l2-stats \(Flat-File Accounting Options\) on page 646](#)
- [mib-profile on page 647](#)
- [mpls \(Security Forwarding Options\) on page 648](#)
- [nonpersistent on page 649](#)
- [object-names on page 650](#)
- [operation on page 651](#)
- [overall-packet \(Flat-File Accounting Options\) on page 652](#)
- [push-backup-to-master \(Accounting Options\) on page 653](#)
- [routing-engine-profile on page 654](#)
- [schema-version \(Flat-File Accounting Options\) on page 655](#)
- [size on page 656](#)
- [source-classes on page 657](#)
- [start-time on page 658](#)
- [traceoptions \(System Accounting\) on page 659](#)
- [transfer-interval on page 661](#)

accounting-options

Supported Platforms M Series, MX Series, PTX Series, SRX Series, T Series, vSRX

Syntax

```

accounting-options {
  class-usage-profile profile-name {
    file filename;
    interval minutes;
    destination-classes {
      destination-class-name;
    }
    source-classes {
      source-class-name;
    }
  }
  cleanup-interval days;
  file filename {
    archive-sites {
      site-name;
    }
    backup-on-failure (master-and-slave | master-only);
    compress;
    files number;
    nonpersistent;
    push-backup-to-master;
    size bytes;
    start-time time;
    transfer-interval minutes;
  }
  filter-profile profile-name {
    counters {
      counter-name;
    }
    file filename;
    interval minutes;
  }
  flat-file-profile profile-name {
    fields {
      all-fields;
      egress-stats {
        all-fields;
        input-bytes;
        input-packets;
        output-bytes;
        output-packets;
        queue-id;
        red-drop-bytes;
        red-drop-packets;
        tail-drop-packets;
      }
      general-param {
        all-fields;
        accounting-type;
        descr;
      }
    }
  }
}

```

```
    line-id;
    logical-interface;
    nas-port-id;
    physical-interface;
    routing-instance;
    timestamp;
    user-name;
    vlan-id;
  }
  ingress-stats {
    all-fields;
    drop-packets;
    input-bytes;
    input-packets;
    output-bytes;
    output-packets;
    queue-id;
  }
  l2-stats {
    all-fields;
    input-mcast-bytes;
    input-mcast-packets;
  }
  overall-packet {
    all-fields;
    input-bytes;
    input-discards;
    input-errors;
    input-packets;
    inputv6-bytes;
    inputv6-packets;
    output-bytes;
    output-errors;
    output-packets;
    outputv6-bytes;
    outputv6-packets;
  }
  service-accounting;
}
file filename;
format (csv | ipdr)
interval minutes;
schema-version schema-name;
}
interface-profile profile-name {
  allow-clear;
  fields {
    field-name;
  }
  file filename;
  interval minutes;
}
mib-profile profile-name {
  file filename;
  interval seconds;
  object-names {
```

```
        mib-object-name;
    }
    operation operation-name;
}
routing-engine-profile profile-name {
    fields {
        field-name;
    }
    file filename;
    interval minutes;
}
}
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure options for accounting statistics collection.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level on page 295](#)
- [Accounting Options Configuration on page 296](#)

allow-clear (Accounting Options)

Supported Platforms [MX Series](#)

Syntax allow-clear;

Hierarchy Level [\[edit accounting-options interface-profile profile-name\]](#)

Release Information Statement introduced in Junos OS Release 17.3R1 on MX Series routers.

Description Reports accounting statistics values cleared by the **clear interfaces statistics** command on a logical interface configured with accounting options to the accounting options flat file associated with the interface. By default, this statement is disabled and the cleared statistics are not reported to the flat file, although the counters do show as cleared in the CLI.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Interface Profile on page 309](#)
- [Accounting Options Configuration on page 296](#)

archive-sites

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax

```
archive-sites {  
    site-name;  
}
```

Hierarchy Level [edit accounting-options [file](#) *filename*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure an archive site. If more than one site name is configured, an ordered list of archive sites for the accounting-data log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with a filename of the format *router-name_log-filename_timestamp*.

Options *site-name*—Any valid FTP/SCP URL to a destination.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring Archive Sites for the Files on page 306](#)

backup-on-failure (Accounting Options)

Supported Platforms [MX Series](#)

Syntax backup-on-failure (master-and-slave | master-only);

Hierarchy Level [edit accounting-options [file](#) *filename*]

Release Information Statement introduced in Junos OS Release 16.1.

Description Configure the router to save a copy of the accounting file locally, to the `/var/log/pfedBackup` directory of the relevant Routing Engine, in the event that file transfer to the remote archive sites cannot be completed.

Options **master-and-slave**—Back up accounting files from both the master Routing Engine and the backup Routing Engine.

master-only—Back up accounting files from only the master Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Accounting-Data Log Files on page 304](#)
- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

class-usage-profile

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax `class-usage-profile profile-name {
 file filename;
 interval minutes;
 source-classes {
 source-class-name;
 }
 destination-classes {
 destination-class-name;
 }
}`

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Create a class usage profile, which is used to log class usage statistics to a file in the `/var/log` directory. The class usage profile logs class usage statistics for the configured source classes on every interface that has **destination-class-usage** configured.

For information about configuring source classes, see the [Junos Routing Protocols Configuration Guide](#). For information about configuring source class usage, see the [Junos Network Management Configuration Guide](#).

Options *profile-name*—Name of the destination class profile.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Class Usage Profiles on page 320](#)

compress (Accounting Options)

Supported Platforms	MX Series
Syntax	compress;
Hierarchy Level	[edit accounting-options file filename]
Release Information	Statement introduced in Junos OS Release 16.1.
Description	Compress the accounting file during file transfer to the backup site.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Accounting-Data Log Files on page 304• Configuring Flat-File Accounting for Layer 2 Wholesale• Configuring Flat-File Accounting for Extensible Subscriber Services Management• Flat-File Accounting Overview

counters

Supported Platforms	EX Series , M Series , MX Series , PTX Series , SRX Series , T Series , vSRX
Syntax	<pre>counters { counter-name; }</pre>
Hierarchy Level	[edit accounting-options filter-profile profile-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Names of counters for which filter profile statistics are collected. The packet and byte counts for the counters are logged to a file in the <code>/var/log</code> directory.
Options	<i>counter-name</i> —Name of the counter.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Counters on page 313

destination-classes

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<pre>destination-classes { destination-class-name; }</pre>
Hierarchy Level	[edit accounting-options class-usage-profile <i>profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 15.1F6 for PTX Series routers with third-generation FPCs installed.
Description	Specify the destination classes for which statistics are collected.
Options	<i>destination-class-name</i> —Name of the destination class to include in the source class usage profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Class Usage Profile on page 320

egress-stats (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax

```
egress-stats {  
    all-fields;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
    red-drop-bytes;  
    red-drop-packets;  
    tail-drop-packets;  
}
```

Hierarchy Level [edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify egress queue statistics to be collected for the interface.

Options **all-fields**—Collect all egress queue statistics available for the interface context, logical or physical.

input-bytes—Collect the number of octets queued including traffic dropped because of congestion.

input-packets—Collect the number of packets queued including traffic dropped because of congestion.

output-bytes—Collect the number of octets transmitted by the egress queue.

output-packets—Collect the number of packets transmitted by the egress queue.

queue-id—Collect the logical identifier for the egress queue; identifies the traffic class.

red-drop-bytes—Collect the number of octets dropped on the egress queue because of random early detection.

red-drop-packets—Collect the number of packets dropped on the egress queue because of random early detection.

tail-drop-packets—Collect the number of packets dropped in the egress queue because of tail drop.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Flat-File Accounting for Layer 2 Wholesale*
 - *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
 - *Flat-File Accounting Overview*

fields (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

```
Syntax  fields {
        all-fields;
        egress-stats {
            all-fields;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
            red-drop-bytes;
            red-drop-packets;
            tail-drop-packets;
        }
        general-param {
            all-fields;
            accounting-type;
            descr;
            line-id;
            logical-interface;
            nas-port-id;
            physical-interface;
            routing-instance;
            timestamp;
            user-name;
            vlan-id;
        }
        ingress-stats {
            all-fields;
            drop-packets;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
        }
        l2-stats {
            all-fields;
            input-mcast-bytes;
            input-mcast-packets;
        }
        overall-packet {
            all-fields;
            input-bytes;
            input-discards;
            input-errors;
            input-packets;
            inputv6-bytes;
            inputv6-packets;
            output-bytes;
            output-errors;
```

```
    output-packets;  
    outputv6-bytes;  
    outputv6-packets;  
  }  
  service-accounting;  
}
```

Hierarchy Level [edit accounting-options [flat-file-profile](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify the accounting statistics and the nonstatistical information to be collected for an interface and recorded in the accounting flat file created by the profile.

Options **all-fields**—Include all available statistical fields in the accounting file. Many fields are available for both logical interfaces and physical interfaces, but some fields are available only for one or the other interface type.

service-accounting—Include the filter counts in bytes for the inet input filter, inet output filter, inet6 input filter, and inet6 output filter in the service accounting flat file. Statistics reported are the running total values.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Configuring Service Accounting in Local Flat Files*
- *Flat-File Accounting Overview*

fields (for Interface Profiles)

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax

```
fields {  
    field-name;  
}
```

Hierarchy Level [edit accounting-options [interface-profile](#) *profile-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Statistics to collect in an accounting-data log file for an interface.

Options *field-name*—Name of the field:

- **input-bytes**—Input bytes
- **input-errors**—Generic input error packets
- **input-multicast**—Input packets arriving by multicast
- **input-packets**—Input packets
- **input-unicast**—Input unicast packets
- **output-bytes**—Output bytes
- **output-errors**—Generic output error packets
- **output-multicast**—Output packets sent by multicast
- **output-packets**—Output packets
- **output-unicast**—Output unicast packets

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Interface Profile on page 309](#)

fields (for Routing Engine Profiles)

Supported Platforms	M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<pre>fields { field-name; }</pre>
Hierarchy Level	[edit accounting-options routing-engine-profile <i>profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Statistics to collect in an accounting-data log file for a Routing Engine.
Options	<p><i>field-name</i>—Name of the field:</p> <ul style="list-style-type: none">• cpu-load-1—Average system load over the last 1 minute• cpu-load-5—Average system load over the last 5 minutes• cpu-load-15—Average system load over the last 15 minutes• date—Date, in YYYYMMDD format• host-name—Hostname for the router• time-of-day—Time of day, in HHMMSS format• uptime—Time since last reboot, in seconds
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Routing Engine Profile on page 325

file (Associating with a Profile)

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	file <i>filename</i> ;
Hierarchy Level	[edit accounting-options class-usage-profile profile-name], [edit accounting-options filter-profile profile-name], [edit accounting-options interface-profile profile-name], [edit accounting-options mib-profile profile-name], [edit accounting-options routing-engine-profile profile-name]
Release Information	Statement introduced before Junos OS Release 7.4. The [edit accounting-options mib-profile profile-name] hierarchy added in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series Switches.
Description	Specify the accounting log file associated with the profile.
Options	<i>filename</i> —Name of the log file. You must specify a filename already configured in the file statement at the [edit accounting-options] hierarchy level.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Profile on page 309• Configuring the Filter Profile on page 312• Configuring the MIB Profile on page 322• Configuring the Routing Engine Profile on page 325

file (Configuring a Log File)

Supported Platforms EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX

Syntax

```
file filename {  
    archive-sites {  
        site-name;  
    }  
    backup-on-failure (master-and-slave | master-only);  
    compress;  
    files number;  
    nonpersistent;  
    push-backup-to-master;  
    size bytes;  
    start-time time;  
    transfer-interval minutes;  
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify a log file to be used for accounting data.

Options *filename*—Name of the file in which to write accounting data.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Accounting-Data Log Files on page 304](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale](#)
- [Configuring Flat-File Accounting for Extensible Subscriber Services Management](#)
- [Flat-File Accounting Overview](#)

file (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax `file filename;`

Hierarchy Level [edit accounting-options [flat-file-profile](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify the name of the accounting file created by a flat-file profile. By default, the filename becomes the name of the local directory where the accounting file is backed up: `/var/log/pfedBackup/filename`.

Options *filename*—Name of the accounting file. The complete output filename is in the format *filename.hostname.file-number_timestamp.gz*.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Flat-File Accounting for Layer 2 Wholesale](#)
- [Configuring Flat-File Accounting for Extensible Subscriber Services Management](#)
- [Flat-File Accounting Overview](#)

files

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	files <i>number</i> ;
Hierarchy Level	[edit accounting-options file <i>filename</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the maximum number of log files to be used for accounting data.
Options	<i>number</i> —The maximum number of files. When a log file (for example, profilelog) reaches its maximum size, it is renamed profilelog.0 , then profilelog.1 , and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. The minimum value for <i>number</i> is 3 and the default value is 10.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Accounting-Data Log Files on page 304

filter-profile

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax `filter-profile profile-name {
 counters {
 counter-name;
 }
 file filename;
 interval minutes;
}`

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Create a profile to filter and collect packet and byte count statistics and write them to a file in the `/var/log` directory. To apply the profile to a firewall filter, you include the **accounting-profile** statement at the `[edit firewall filter filter-name]` hierarchy level. For more information about firewall filters, see *Firewall Filters Overview*.

Options *profile-name*—Name of the filter profile.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Firewall Filters Overview](#)
- [Configuring the Filter Profile on page 312](#)

flat-file-profile (Accounting Options)

Supported Platforms [MX Series](#)

```
Syntax flat-file-profile profile-name{
    fields {
        all-fields;
        egress-stats {
            all-fields;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
            red-drop-bytes;
            red-drop-packets;
            tail-drop-packets;
        }
        general-param {
            all-fields;
            accounting-type;
            descr;
            line-id;
            logical-interface;
            nas-port-id;
            physical-interface;
            routing-instance;
            timestamp;
            user-name;
            vlan-id;
        }
        ingress-stats {
            all-fields;
            drop-packets;
            input-bytes;
            input-packets;
            output-bytes;
            output-packets;
            queue-id;
        }
        l2-stats {
            all-fields;
            input-mcast-bytes;
            input-mcast-packets;
        }
        overall-packet {
            all-fields;
            input-bytes;
            input-discards;
            input-errors;
            input-packets;
            inputv6-bytes;
            inputv6-packets;
            output-bytes;
```

```
        output-errors;  
        output-packets;  
        outputv6-bytes;  
        outputv6-packets;  
    }  
    service-accounting;  
}  
file filename;  
format (csv | ipdr)  
interval minutes;  
schema-version schema-name;  
}
```

Hierarchy Level [edit [accounting-options](#)]

Release Information Statement introduced in Junos OS Release 16.1R4.
service-accounting option added in Junos OS Release 17.1.

Description Configure a flat-file accounting profile that defines the contents of a flat file that records accounting statistics collected from the Packet Forwarding Engine for an interface at regular intervals. To be used, the profile is associated with a subscriber interface. The accounting flat file is archived by the accounting-options archiving mechanism.

Options *profile-name*—Name of the flat-file profile.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- [Configuring Accounting-Data Log Files on page 304](#)
- [Configuring Flat-File Accounting for Layer 2 Wholesale](#)
- [Configuring Flat-File Accounting for Extensible Subscriber Services Management](#)
- [Configuring Service Accounting in Local Flat Files](#)
- [Flat-File Accounting Overview](#)
- [show extensible-subscriber-services accounting](#)

format (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax `format (csv | ipdr);`

Hierarchy Level [edit accounting-options [flat-file-profile](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify the format for logging the flat-file accounting statistics.

Options **csv**—Comma-separated values (CSV) format.

ipdr—IP Detail Record (IPDR) format.

Default: ipdr

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

general-param (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax `general-param {
 all-fields;
 accounting-type;
 descr;
 line-id;
 logical-interface;
 nas-port-id;
 physical-interface;
 routing-instance;
 timestamp;
 user-name;
 vlan-id;
}`

Hierarchy Level [edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information Statement introduced in Junos OS Release 16.1R4.
 user-name option added in Junos OS Release 17.1.

Description Specify general, nonstatistical interface parameters that are displayed as part of the header for the accounting file.

Options **all-fields**—Display all available nonstatistical fields. Many fields are available for both logical interfaces and physical interfaces, but some fields are available for only one interface type.

accounting-type—(Logical interfaces only) Display the accounting status type.

descr—Display the description of the interface as configured.

line-id—(Logical interfaces only) Display the access line identifier.

logical-interface—(Logical interfaces only) Display the name of the logical interface.

nas-port-id—(Logical interfaces only) Display the NAS port ID.

physical-interface—(Physical interfaces only) Display the name of the physical interface.

routing-instance—Display the name of the routing instance to which the interface belongs.

timestamp—Display the timestamp of the accounting record.

user-name—Display the name of the subscriber.

vlan-id—(Logical interfaces only) Display the VLAN identifier.

Required Privilege	system—To view this statement in the configuration.
Level	system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Flat-File Accounting for Layer 2 Wholesale</i>• <i>Configuring Flat-File Accounting for Extensible Subscriber Services Management</i>• <i>Configuring Service Accounting in Local Flat Files</i>• <i>Flat-File Accounting Overview</i>

ingress-stats (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax

```
ingress-stats {  
    all-fields;  
    drop-packets;  
    input-bytes;  
    input-packets;  
    output-bytes;  
    output-packets;  
    queue-id;  
}
```

Hierarchy Level [edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify ingress queue statistics to be collected for the interface.

Options **all-fields**—Collect all ingress queue statistics available for the interface context, logical or physical.

drop-packets—Collect the number of received packets dropped on the Ingress queue.

input-bytes—Collect the number of octets received on the queue for the traffic class indicated by the queue identifier.

input-packets—Collect the number of packets received on the queue for the traffic class indicated by the queue identifier.

output-bytes—Collect the number of octets forwarded for the traffic class indicated by the queue identifier. Same value as **input-bytes** unless oversubscription is present at the ingress.

output-packets—Collect the number of packets forwarded for the traffic class indicated by the queue identifier. Same value as **input-packets** unless oversubscription is present at the ingress.

queue-id—Collect the logical identifier for the ingress queue; identifies the traffic class.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

interface-profile

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax `interface-profile profile-name {
 allow-clear;
 fields {
 field-name;
 }
 file filename;
 interval minutes;
}`

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 on EX Series switches.

Description Create a profile to filter and collect error and packet statistics and write them to a file in the `/var/log` directory. You can specify an interface profile for either a physical or a logical interface.

Options *profile-name*—Name of the interface profile.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Interface Profile on page 309](#)
- [Accounting Options Configuration on page 296](#)

interval

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	interval <i>minutes</i> ;
Hierarchy Level	[edit accounting-options class-usage-profile profile-name], [edit accounting-options filter-profile profile-name], [edit accounting-options interface-profile profile-name], [edit accounting-options mib-profile profile-name], [edit accounting-options routing-engine-profile profile-name]
Release Information	Statement introduced before Junos OS Release 7.4. The [edit accounting-options mib-profile profile-name] hierarchy level added in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify how often statistics are collected for the accounting profile.
Options	minutes —Length of time between each collection of statistics. Range: 1 through 2880 minutes Default: 30 minutes
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interface Profile on page 309• Configuring the Filter Profile on page 312• Configuring the MIB Profile on page 322• Configuring the Routing Engine Profile on page 325

interval (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax `interval minutes;`

Hierarchy Level `[edit accounting-options flat-file-profile profile-name]`

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify the interval in minutes at which the Packet Forwarding Engine associated with the interface is polled to collect the statistics specified in the flat-file accounting profile. These interim accounting results are recorded in the flat file.



NOTE: The value configured with this statement is superseded by the value configured with the `update-interval` statement at the `[edit access profile profile-name service accounting]` hierarchy level. That access profile interval value is in turn superseded by an update interval value configured in the RADIUS attribute, Service-Interim-Acct-Interval (VSA 26–140).

Options *minutes*—Polling interval.
Range: 1 through 2880 minutes
Default: 15 minutes

Required Privilege Level `system`—To view this statement in the configuration.
`system-control`—To add this statement to the configuration.

Related Documentation

- [Configuring Flat-File Accounting for Layer 2 Wholesale](#)
- [Configuring Flat-File Accounting for Extensible Subscriber Services Management](#)
- [Configuring Service Accounting in Local Flat Files](#)
- [Flat-File Accounting Overview](#)

l2-stats (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax

```
l2-stats {  
    all-fields;  
    input-mcast-bytes;  
    input-mcast-packets;  
}
```

Hierarchy Level [edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify the statistics to collect for the named flat-file-profile field.

Options **all-fields**—Collect all Layer 2 statistics for the named flat-file profile.

input-mcast-bytes—Collect multicast bytes from the input side for the named flat-file profile.

input-mcast-packets—Collect multicast packets from the input side for the named flat-file profile.

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

Related Documentation

- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

mib-profile

Supported Platforms EX Series, M Series, MX Series, SRX Series, T Series, vSRX

Syntax

```
mib-profile profile-name {
  file filename;
  interval minutes;
  object-names {
    mib-object-name;
  }
  operation operation-name;
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Create a MIB profile to collect selected MIB statistics and write them to a file in the `/var/log` directory.



NOTE: Do not configure MIB objects related to interface octets or packets for a MIB profile, because it can cause the SNMP walk or a CLI show command to time out.

Options *profile-name*—Name of the MIB statistics profile.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the MIB Profile on page 322](#)

mpls (Security Forwarding Options)

Supported Platforms [SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX](#)

Syntax

```
mpls {  
    mode packet-based;  
}
```

Hierarchy Level [edit security forwarding-options family]

Release Information Statement introduced in Junos OS Release 9.0.

Description Enable the forwarding of MPLS traffic. By default, the device drops MPLS traffic.



CAUTION: Because MPLS operates in packet mode, security services are not available.



NOTE: Packet-based processing is not supported on the following SRX Series devices: SRX1500, SRX5600, and SRX5800.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [MPLS Overview](#)

nonpersistent

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax nonpersistent;

Hierarchy Level [edit accounting-options [file filename](#)]

Release Information Statement introduced in Junos OS Release 8.3.

Description Store log files used for accounting data in the **mfs/var/log** directory (located on DRAM) instead of the **cf/var/log** directory (located on the compact flash drive). This feature is useful for minimizing read/write traffic on the router's compact flash drive.



NOTE: If log files for accounting data are stored on DRAM, these files are lost when you reboot the router. Therefore, you should back up these files periodically.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Storage Location of the File on page 307](#)

object-names

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	<pre>object-names { mib-object-name; }</pre>
Hierarchy Level	[edit accounting-options mib-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the name of each MIB object for which MIB statistics are collected for an accounting-data log file.
Options	<i>mib-object-name</i> —Name of a MIB object. You can specify more than one MIB object name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the MIB Profile on page 322

operation

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<code>operation operation-name;</code>
Hierarchy Level	[edit accounting-options mib-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the name of the operation used to collect MIB statistics for an accounting-data log file.
Options	<i>operation-name</i> —Name of the operation to use. You can specify a get , get-next , or walk operation. Default: walk
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the MIB Profile on page 322

overall-packet (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax `overall-packet {
 all-fields;
 input-bytes;
 input-discards;
 input-errors;
 input-packets;
 inputv6-bytes;
 inputv6-packets;
 output-bytes;
 output-errors;
 output-packets;
 outputv6-bytes;
 outputv6-packets;
}`

Hierarchy Level [edit accounting-options flat-file-profile *profile-name* [fields](#)]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify overall packet statistics to be collected for the interface.

Options **all-fields**—Collect all overall packet statistics available for the interface context, logical or physical.

input-bytes—Collect the number of octets received on the interface.

input-discards—(Physical interfaces only) Collect the number of received packets that were discarded on the interface.

input-errors—(Physical interfaces only) Collect the number of frames with errors received on the interface.

input-packets—Collect the number of packets received on the interface.

input-v6-bytes—Collect the number of IPv6 octets received on the interface.

input-v6-packets—Collect the number of IPv6 packets received on the interface.

output-bytes—Collect the number of octets transmitted on the interface.

output-errors—(Physical interfaces only) Collect the number of frames that could not be transmitted on the interface because of errors.


output-packets—Collect the number of packets transmitted on the interface.

output-v6-bytes—Collect the number of IPv6 octets transmitted on the interface.

output-v6-packets—Collect the number of IPv6 packets transmitted on the interface.

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Flat-File Accounting for Layer 2 Wholesale• Configuring Flat-File Accounting for Extensible Subscriber Services Management• Flat-File Accounting Overview

push-backup-to-master (Accounting Options)

Supported Platforms	MX Series
Syntax	push-backup-to-master;
Hierarchy Level	[edit accounting-options file <i>filename</i>]
Release Information	Statement introduced in Junos OS Release 16.1R4.
Description	<p>Configure the router to save the accounting files from the new backup Routing Engine to the new master Routing Engine when a change in mastership occurs. The files are saved to the <code>/var/log/pfedBackup</code> directory on the router. The master Routing Engine includes these accounting files with its own current accounting files when it transfers the files from the backup directory to the archive site at the next transfer interval. Use this statement when the new backup Routing Engine is not able to connect to the archive site; for example, when site is not connected by means of an out-of-band interface or the path to the site is routed through a line card.</p>
	<div>NOTE: The backup Routing Engine's files on the master Routing Engine are sent at each interval even though the files remain the same. If this is more activity than you want, consider using the <code>backup-on-failure master-and-slave</code> statement instead.</div>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Accounting-Data Log Files on page 304• Configuring Flat-File Accounting for Layer 2 Wholesale• Configuring Flat-File Accounting for Extensible Subscriber Services Management• Flat-File Accounting Overview

routing-engine-profile

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax `routing-engine-profile profile-name {
 fields {
 field-name;
 }
 file filename;
 interval minutes;
}`

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Create a Routing Engine profile to collect selected Routing Engine statistics and write them to a file in the `/var/log` directory.

Options *profile-name*—Name of the Routing Engine statistics profile.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Routing Engine Profile on page 325](#)

schema-version (Flat-File Accounting Options)

Supported Platforms [MX Series](#)

Syntax `schema-version schema-name;`

Hierarchy Level [edit accounting-options [flat-file-profile](#) *profile-name*]

Release Information Statement introduced in Junos OS Release 16.1R4.

Description Specify the name of the XML schema that defines the contents and format of the accounting file, and appears in the accounting record header.

Options ***schema-name***—Name of the schema.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Configuring Flat-File Accounting for Layer 2 Wholesale*
- *Configuring Flat-File Accounting for Extensible Subscriber Services Management*
- *Flat-File Accounting Overview*

size

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax `size bytes;`

Hierarchy Level [edit accounting-options [file](#) *filename*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify attributes of an accounting-data log file.

Options **bytes**—Maximum size of each log file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a log file (for example, **profilelog**) reaches its maximum size, it is renamed **profilelog.0**, then **profilelog.1**, and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. If you do not specify a size, the file is closed, archived, and renamed when the time specified for the transfer interval is exceeded.

Syntax: *x* to specify bytes, *xk* to specify KB, *xm* to specify MB, *xg* to specify GB

Range: 256 KB through 1 GB

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Maximum Size of the File on page 305](#)

source-classes

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<pre>source-classes { source-class-name; }</pre>
Hierarchy Level	[edit accounting-options class-usage-profile <i>profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 15.1F6 for PTX Series routers with third-generation FPCs installed.
Description	Specify the source classes for which statistics are collected.
Options	<i>source-class-name</i> —Name of the source class to include in the class usage profile.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a Class Usage Profile on page 320

start-time

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	start-time <i>time</i> ;
Hierarchy Level	[edit accounting-options file filename]
Release Information	Statement introduced in Junos OS Release 8.2. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the start time for transfer of an accounting-data log file.
Options	<i>time</i> —Start time for file transfer. Syntax: <i>YYYY-MM-DD.hh:mm</i>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Start Time for File Transfer on page 308

traceoptions (System Accounting)

Supported Platforms EX Series, MX Series

Syntax `traceoptions {
 file filename <files number> <size size> <world-readable | no-world-readable>;
 flag (all| config | events | radius | tacplus);
 no-remote-trace
}`

Hierarchy Level [edit system accounting]]

Release Information Statement introduced in Junos OS Release 14.2.
tacplus option introduced in Junos OS Release 15.1.

Description Define tracing operations for System Accounting.

Default Trace options are not enabled by default.

Options **file *filename***—Name of the file in which Junos OS stores the accounting logs. By default, this is created under the /var/log directory.

files *number*—(Optional) Maximum number of trace files. When a trace file reaches the size specified by the size option, the filename is appended with 0 and compressed. For example, when trace file named trace-file-log reaches size <*size*>, it is renamed and compressed to trace-file-log.0.gz. When trace-file-log reaches size <*size*> or the second time, the trace-file-log.0.gz is renamed to trace-file-log.1.gz and trace-file-log is renamed and compressed to trace-file-log.0.gz. This renaming scheme ensures that the older logs to have a greater index number. When number of trace files reach <*number*> then the oldest file is deleted.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files

Default: 10

flag *flag*—Tracing operation to perform. You can include one or more of the following flags:

- **all**—Trace all operations.
- **config**—Trace configuration processing.
- **events**—Trace accounting events and their processing.
- **radius**—Trace RADIUS processing.
- **tacplus**—Trace TACPLUS processing.

no-remote-trace—(Optional) Disable tracing and logging operations that track normal operations, error conditions, and packets that are generated by or passed through the Juniper Networks device.

no-world-readable—Restrict access to the trace files to the owner.

Default: no-world-readable

size *size*—(Optional) Maximum size of each trace file in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). If you do not specify a unit, the default is bytes. If you specify a maximum file size, you also must specify a maximum number of trace files by using the **files** option and a filename by using the **file** option.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: *size* to specify bytes, *sizek* to specify KB, *sizem* to specify MB, or *sizeg* to specify GB.

Range: 10 KB through 1 MB

Default: 128 KB

world-readable—Enable any user to access the trace files.

Required Privilege	admin—To view this statement in the configuration.
Level	admin-control—To add this statement to the configuration.

transfer-interval

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	transfer-interval <i>minutes</i> ;
Hierarchy Level	[edit accounting-options file filename]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the length of time the file remains open and receives new statistics before it is closed and transferred to an archive site.
Options	minutes —Time the file remains open and receives new statistics before it is closed and transferred to an archive site. Range: 5 through 2880 minutes Default: 30 minutes
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Transfer Interval of the File on page 308

CHAPTER 36

Configuration Statements: Chassis Cluster

- [cluster \(Chassis\) on page 664](#)
- [global-threshold on page 665](#)
- [global-weight on page 666](#)
- [ip-monitoring on page 667](#)
- [ip-monitoring \(Services\) on page 669](#)
- [next-hop on page 670](#)

cluster (Chassis)

Supported Platforms [SRX Series, vSRX](#)

```
Syntax cluster {
    configuration-synchronize {
        no-secondary-bootup-auto;
    }
    control-link-recovery;
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    network-management {
        cluster-master;
    }
    redundancy-group group-number {
        gratuitous-arp-count number;
        hold-down-interval number;
        interface-monitor interface-name {
            weight number;
        }
    }
    ip-monitoring {
        family {
            inet {
                ipv4-address {
                    interface {
                        logical-interface-name;
                        secondary-ip-address ip-address;
                    }
                    weight number;
                }
            }
        }
        global-threshold number;
        global-weight number;
        retry-count number;
        retry-interval seconds;
    }
    node (0 | 1 ) {
        priority number;
    }
    preempt;
}
reth-count number;
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        (world-readable | no-world-readable);
        size maximum-file-size;
    }
    flag flag;
    level {
        (alert | all | critical | debug | emergency | error | info | notice | warning);
```

```

    }
    no-remote-trace;
  }
}

```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 9.0.

Description Configure a chassis cluster.

Options The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [ip-monitoring on page 667](#)

[global-threshold](#)

Supported Platforms [SRX Series, vSRX](#)

Syntax global-threshold *number*;

Hierarchy Level [edit chassis cluster redundancy-group *group-number* ip-monitoring]

Release Information Statement introduced in Junos OS Release 10.1.

Description Specify the failover value for all IP addresses monitored by the redundancy group. When IP addresses with a configured total weight in excess of the threshold have become unreachable, the weight of IP monitoring is deducted from the redundancy group threshold.

Options *number* —Value at which the IP monitoring weight is applied against the redundancy group failover threshold.

Range: 0 through 255

Default: 0

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation • [ip-monitoring on page 667](#)

global-weight

Supported Platforms [SRX Series, vSRX](#)

Syntax `global-weight number;`

Hierarchy Level `[edit chassis cluster redundancy-group group-number ip-monitoring]`

Release Information Statement introduced in Junos OS Release 10.1.

Description Specify the relative importance of all IP address monitored objects to the operation of the redundancy group. Every monitored IP address is assigned a weight. If the monitored address becomes unreachable, the weight of the object is deducted from the global-threshold of IP monitoring objects in its redundancy group. When the global-threshold reaches 0, the global-weight is deducted from the redundancy group. Every redundancy group has a default threshold of 255. If the threshold reaches 0, a failover is triggered. Failover is triggered even if the redundancy group is in manual failover mode and preemption is not enabled.

Options *number*—Combined weight assigned to all monitored IP addresses. A higher weight value indicates a greater importance.

Range: 0 through 255

Default: 255

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation

- [ip-monitoring on page 667](#)

ip-monitoring

Supported Platforms SRX Series, vSRX

Syntax

```
ip-monitoring {
  family {
    inet {
      ipv4-address {
        interface {
          logical-interface-name;
          secondary-ip-address ip-address;
        }
        weight number;
      }
    }
  }
  global-threshold number;
  global-weight number;
  retry-count number;
  retry-interval seconds;
}
```

Hierarchy Level [edit chassis cluster redundancy-group *group-number*]

Release Information Statement updated in Junos OS Release 10.1.

Description Specify a global IP address monitoring threshold and weight, and the interval between pings (**retry-interval**) and the number of consecutive ping failures (**retry-count**) permitted before an IP address is considered unreachable for all IP addresses monitored by the redundancy group. Also specify IP addresses, a monitoring weight, a redundant Ethernet interface number, and a secondary IP monitoring ping source for each IP address, for the redundancy group to monitor.

Options *IPv4 address*—The address to be continually monitored for reachability.



NOTE: All monitored object failures, including IP monitoring, are deducted from the redundancy group threshold priority. Other monitored objects include interface monitor, SPU monitor, cold-sync monitor, and NPC monitor (on supported platforms).

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface	To view this statement in the configuration.
interface-control	To add this statement to the configuration.

**Related
Documentation**

- [interface \(Chassis Cluster\)](#)
- [global-threshold on page 665](#)
- [global-weight on page 666](#)
- [weight](#)
- [Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring on page 382](#)

ip-monitoring (Services)

Supported Platforms [SRX Series, vSRX](#)

Syntax

```
ip-monitoring {
  policy policy-name {
    match {
      rpm-probe [probe-name];
    }
    no-preempt;
    then {
      interface interface-name (disable | enable);
      preferred-route {
        route destination-address {
          discard
          next hop next-hop;
          preferred-metric metric;
        }
        routing-instances name;
      }
    }
  }
}
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 10.4.
Support added for the **discard** option starting from Junos OS Release 15.1X49-D60.

Description Configure IP monitoring.

Options The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level services—To view this statement in the configuration.
services-control—To add this statement to the configuration.

Related Documentation

- [icmp on page 679](#)

next-hop

Supported Platforms [vSRX](#)

Syntax `next-hop next-hop;`

Hierarchy Level `[edit services rpm probe owner test test-name]`

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the next-hop address to which the probe should be sent.

Required Privilege Level `services`—To view this statement in the configuration.
`services-control`—To add this statement to the configuration.

Related Documentation

- [probe](#)

CHAPTER 37

Configuration Statements: Datapath Debug

- [action-profile](#) on page 672
- [capture-file \(Security\)](#) on page 674
- [datapath-debug](#) on page 675
- [flow \(Security Flow\)](#) on page 677
- [icmp](#) on page 679
- [maximum-capture-size \(Datapath Debug\)](#) on page 679
- [traceoptions \(Security Datapath Debug\)](#) on page 680

action-profile

Supported Platforms SRX5400, SRX5600, SRX5800, vSRX

Syntax `action-profile profile-name {
 event (jexec | lbt | lt-enter | lt-leave | mac-egress | mac-ingress | np-egress | np-ingress |
 pot) {
 count;
 packet-dump;
 packet-summary;
 trace;
 }
 module {
 flow {
 flag {
 all;
 }
 }
 }
 preserve-trace-order;
 record-pic-history;
}`

Hierarchy Level [edit security datapath-debug]

Release Information Command introduced in Junos OS Release 10.0.

Description Configure the action profile options for data path debugging.

- Options**
- ***action-profile name*** — Name of the action profile.
 - **event**—Enable the events to trace the packet when the packet hit the events (jexec, lbt, lt-enter, lt-leave, mac-egress, mac-ingress, np-egress, np-ingress, pot)
 - **count**—Number of times a packet hits the specified event.
 - **packet-dump**—Capture the packet that hits the specified event.
 - **packet-summary**—Print the source/destination IP address details with protocol number and IP length details along with trace message for the specified event.
 - **trace**—Print the standard trace message when the packet hits the specified event.
 - **module**—Turn on the flow session related trace messages.
 - **flow**—Trace flow session related messages.
 - **flag**—Specify which flow message needs to be traced.
 - **all**—Trace all possible flow trace messages.
 - **trace**—Print the standard trace message when the packet hits the specified event.
 - **preserve-trace-order**—Preserve trace order.

- **record-pic-history**—Record the PICs in which the packet has been processed.

Required Privilege Level security—To view this statement in the configuration.
 security-control—To add this statement to the configuration.

Related Documentation • [Example: Configuring Packet Capture for Datapath Debugging on page 575](#)

capture-file (Security)

Supported Platforms [SRX1500, SRX5400, SRX5600, SRX5800, vSRX](#)

Syntax `capture-file {
 filename;
 files number;
 format pcap-format;
 size maximum-file-size;
 (world-readable | no-world-readable);
}`

Hierarchy Level [edit security datapath-debug]

Release Information Statement introduced in Junos OS Release 10.4.

Description Sets packet capture for performing the datapath-debug action.

- Options**
- **filename**—Name of the file to receive the output of the packet capturing operation.
 - **files**—Maximum number of capture files.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 1 through 10 files
 - **format**—Describes the format of the capture file. The default format file is pcap. You can also set it as private (binary) format.
 - **size**—Describes the size limit of the capture file.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Range: 10 KB through 100 MB
 - **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [System Log Messages](#)

datapath-debug

Supported Platforms [SRX Series, vSRX](#)

```
Syntax datapath-debug {
    action-profile profile-name {
        event (jexec | lbt | lt-enter | lt-leave | mac-egress | mac-ingress | np-egress | np-ingress
            | pot) {
            count;
            packet-dump;
            packet-summary;
            trace;
        }
        module {
            flow {
                flag {
                    all;
                }
            }
        }
        preserve-trace-order;
        record-pic-history;
    }
    capture-file {
        filename;
        files number;
        format pacp-format;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    maximum-capture-size value;
    packet-filter packet-filter-name {
        action-profile (profile-name | default);
        destination-port (port-range | protocol-name);
        destination-prefix destination-prefix;
        interface logical-interface-name;
        protocol (protocol-number | protocol-name);
        source-port (port-range | protocol-name);
        source-prefix source-prefix;
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        no-remote-trace;
    }
}
```

Hierarchy Level [\[edit security\]](#)

Release Information Command introduced in Junos OS Release 10.0.

Description Configure the data path debugging options.



NOTE: Data path debugging is not supported on SRX1500 devices.

Options The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege security—To view this statement in the configuration.
Level security-control—To add this statement to the configuration.

Related Documentation

- *Understanding Data Path Debugging for Logical Systems*

flow (Security Flow)

Supported Platforms [SRX Series, vSRX](#)

```
Syntax flow {
    aging {
        early-ageout seconds;
        high-watermark percent;
        low-watermark percent;
    }
    allow-dns-reply;
    ethernet-switching {
        block-non-ip-all;
        bpdv-vlan-flooding;
        bypass-non-ip-unicast;
        no-packet-flooding {
            no-trace-route;
        }
    }
    force-ip-reassembly;
    ipsec-performance-acceleration;
    load distribution {
        session-affinity ipsec;
    }
    packet-log {
        enable;
        throttle-interval;
        packet-filter <filter-name>;
    }
    pending-sess-queue-length (high | moderate | normal);
    route-change-timeout seconds;
    syn-flood-protection-mode (syn-cookie | syn-proxy);
    tcp-mss {
        all-tcp mss value;
        gre-in {
            mss value;
        }
        gre-out {
            mss value;
        }
        ipsec-vpn {
            mss value;
        }
    }
    tcp-session {
        fin-invalidate-session;
        no-sequence-check;
        no-syn-check;
        no-syn-check-in-tunnel;
        rst-invalidate-session;
        rst-sequence-check;
        strict-syn-check;
        tcp-initial-timeout seconds;
        time-wait-state {
```

```
        (session-ageout | session-timeout seconds);
    }
}
traceoptions {
    file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    flag flag;
    no-remote-trace;
    packet-filter filter-name {
        destination-port port-identifier;
        destination-prefix address;
        interface interface-name;
        protocol protocol-identifier;
        source-port port-identifier;
        source-prefix address;
    }
    rate-limit messages-per-second;
}
}
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 9.5.

Description Determine how the device manages packet flow. The device can regulate packet flow in the following ways:

Options The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- *Juniper Networks Devices Processing Overview*
- *Understanding Session Characteristics for SRX Series Services Gateways*
- *Understanding Flow in Logical Systems for SRX Series Devices*

icmp

Supported Platforms	SRX Series, vSRX
Syntax	<pre>icmp{ destination-interface <i>interface-name</i>; }</pre>
Hierarchy Level	[edit services rpm probe-server]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Specify the port information for the ICMP server.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Understanding ICMP Fragment Protection

maximum-capture-size (Datapath Debug)

Supported Platforms	SRX1500, SRX5400, SRX5600, SRX5800
Syntax	maximum-capture-size <i>maximum-capture-size</i> ;
Hierarchy Level	[edit security datapath-debug]
Release Information	Statement introduced in Junos OS Release 10.0.
Description	Specifies maximum packet capture length.
Options	<ul style="list-style-type: none"> • maximum-capture-size <i>maximum-capture-size</i>—Specify the maximum packet capture length. <p>Range: 68 through 10,000 bytes</p>
Required Privilege Level	<p>security—To view this statement in the configuration.</p> <p>security-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • System Log Messages

traceoptions (Security Datapath Debug)

Supported Platforms SRX Series, vSRX

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  no-remote-trace;
}
```

Hierarchy Level [edit security datapath-debug]

Release Information Command introduced in Junos OS Release 9.6.

Description Sets the trace options for datapath-debug.



NOTE: Data path debugging is not supported on SRX1500 devices.

Options

- file**—Configure the trace file options.

- filename**—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, the name of the file is the name of the process being traced.
- files number**—Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed to **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files

Default: 10 files

- match regular-expression**—Refine the output to include lines that contain the regular expression.
- size maximum-file-size**—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the trace-file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme

continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option and a filename.

Syntax: `x K` to specify KB, `x m` to specify MB, or `x g` to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

- **world-readable | no-world-readable**—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option
- **no-remote-trace**—Set remote tracing as disabled.

Required Privilege	<code>trace</code> —To view this statement in the configuration.
Level	<code>trace-control</code> —To add this statement to the configuration.

CHAPTER 38

Configuration Statements: Health Monitoring

- [falling-threshold on page 684](#)
- [health-monitor on page 685](#)
- [idp \(SNMP\) on page 686](#)
- [routing-engine \(SNMP Resource Level\) on page 687](#)
- [interval on page 688](#)
- [rising-threshold on page 689](#)

falling-threshold

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	falling-threshold <i>percentage</i> ;
Hierarchy Level	[edit snmp health-monitor], [edit snmp health-monitor idp]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.3 for the [edit snmp health-monitor idp] hierarchy level.
Description	Set the lower threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —The lower threshold for the alarm entry. Range: 1 through 100 Default: 70 percent of the maximum possible value
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Falling Threshold or Rising Threshold on page 287• rising-threshold on page 689

health-monitor

Supported Platforms EX Series, M Series, MX Series, PTX Series, SRX Series, T Series

Syntax

```
health-monitor {
  falling-threshold percentage;
  interval seconds;
  rising-threshold percentage;
  idp {
    falling-threshold percentage;
    interval seconds;
    rising-threshold percentage;
  }
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Option **idp** introduced in Junos OS Release 9.3 for all supported platforms.

Description Configure health monitoring.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring Health Monitoring on Devices Running Junos OS on page 285](#)

idp (SNMP)

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `idp {
 falling-threshold percentage;
 interval seconds;
 rising-threshold percentage;
}`

Hierarchy Level [edit snmp health-monitor]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure health monitoring for Intrusion Detection and Prevention (IDP).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring Health Monitoring on Devices Running Junos OS on page 285](#)

routing-engine (SNMP Resource Level)

Supported Platforms SRX300, SRX320, SRX340, SRX345, SRX550

Syntax

```
routing-engine {
  resource <cpu | memory | open-files-count | process-count | storage | temperature |
    traceoptions > ;
  {
    interval <interval in secs>;
    moderate-threshold <percentage level>;
    high-threshold <percentage level>;
    critical-threshold <percentage level>;
    action <monitor | prevent | recover>;
  }
}
```

Hierarchy Level [edit snmp health-monitor routing-engine]

Release Information Statement introduced in Junos OS Release 12.1X44-D10. Statement modified in Junos OS Release 15.1X49-D10.

Description Override the global configuration for a resource.

- Options**
- **interval**—Monitoring interval in seconds.
Range: 1 through 604800 seconds
Default: 300 seconds
 - **moderate-threshold**—Percentage of moderate threshold level resource utilization.
Range: 30 through 99 percent
Default: 70 percent
 - **high-threshold** —Percentage of high-threshold level resource utilization.
Range: 30 through 99 percent
Default: 80 percent
 - **critical-threshold** —Percentage of critical threshold level resource utilization.
Range: 30 through 99 percent
Default: 90 percent
 - **action**—Enable action for all resources.
Default: If action is not enabled, the default action is prevent.



WARNING: If the system health management action for an affected resource is configured to recover, then certain intrusive operations necessary for preventing system breakdown are taken. Intrusive operations

can include restarting or terminating processes, deleting files, and so on. Such action information is logged in the system health management history and system log.

Required Privilege	security—To view this statement in the configuration.
Level	security-control—To add this statement to the configuration.

interval

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
----------------------------	--

Syntax	interval <i>seconds</i> ;
---------------	---------------------------

Hierarchy Level	[edit snmp health-monitor], [edit snmp health-monitor idp]
------------------------	---

Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.3 for the [edit snmp health-monitor idp] hierarchy level.
----------------------------	---

Description	Interval between samples.
--------------------	---------------------------

Options	seconds —Time between samples, in seconds. Range: 1 through 2147483647 seconds Default: 300 seconds
----------------	--

Required Privilege	snmp—To view this statement in the configuration.
Level	snmp-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• Configuring the Interval on page 288
------------------------------	--

rising-threshold

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	rising-threshold <i>percentage</i> ;
Hierarchy Level	[edit snmp health-monitor], [edit snmp health-monitor idp]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 9.3 for the [edit snmp health-monitor idp] hierarchy level.
Description	Set the upper threshold for the monitored object when you configure a health monitor alarm. By setting a rising and a falling threshold for a monitored object, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>percentage</i> —The lower threshold for the alarm entry. Range: 1 through 100 Default: 80 percent of the maximum possible value
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• falling-threshold on page 684• Configuring the Falling Threshold or Rising Threshold on page 287

CHAPTER 39

Configuration Statements: Remote Monitoring (RMON)

- [alarm \(SNMP RMON\) on page 692](#)
- [community on page 693](#)
- [description on page 694](#)
- [event on page 695](#)
- [falling-event-index on page 696](#)
- [falling-threshold on page 697](#)
- [falling-threshold-interval on page 698](#)
- [interval on page 699](#)
- [request-type on page 700](#)
- [rising-event-index on page 701](#)
- [rising-threshold on page 702](#)
- [rmon on page 702](#)
- [sample-type on page 703](#)
- [startup-alarm on page 704](#)
- [syslog-subtag on page 705](#)
- [type on page 706](#)
- [variable on page 707](#)

alarm (SNMP RMON)

Supported Platforms ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFX Series, SRX320, T Series

Syntax `alarm index {
 description description;
 falling-event-index index;
 falling-threshold integer;
 falling-threshold-interval seconds;
 interval seconds;
 request-type (get-next-request | get-request | walk-request);
 rising-event-index index;
 rising-threshold integer;
 sample-type (absolute-value | delta-value);
 startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
 syslog-subtag syslog-subtag;
 variable oid-variable;
}`

Hierarchy Level [edit snmp rmon]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure RMON alarm entries.

Options *index*—Identifies this alarm entry as an integer.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring an RMON Alarm Entry and Its Attributes on page 242](#)
- [event on page 695](#)
- *RMON MIB Event, Alarm, Log, and History Control Tables*
- *Monitoring RMON MIB Tables*
- *Understanding RMON*

community

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	community <i>community-name</i> ;
Hierarchy Level	[edit snmp rmon event <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The trap group that is used when generating a trap (if eventType is configured to send traps). If that trap group has the rmon-alarm trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of eventCommunity). If nothing is configured, traps are sent to each group with the rmon-alarm category set.
Options	community-name —Identifies the trap group that is used when generating a trap if the event is configured to send traps.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an RMON Event Entry and Its Attributes on page 246

description

Supported Platforms	EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	<code>description <i>description</i>;</code>
Hierarchy Level	[edit snmp rmon alarm <i>index</i>], [edit snmp rmon event <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Text description of alarm or event.
Options	<i>description</i> —Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Description on page 243• Configuring an RMON Event Entry and Its Attributes on page 246

event

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	<pre>event <i>index</i> { community <i>community-name</i>; description <i>description</i>; type <i>type</i>; }</pre>
Hierarchy Level	[edit snmp rmon]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure RMON event entries.
Options	index —Identifier for a specific event entry. The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an RMON Event Entry and Its Attributes on page 246• alarm (SNMP RMON) on page 692

falling-event-index

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax falling-event-index *index*;

Hierarchy Level [edit snmp rmon [alarm](#) *index*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description The index of the event entry that is used when a falling threshold is crossed. If this value is zero, no event is triggered.

Options *index*—Index of the event entry that is used when a falling threshold is crossed.
Range: 0 through 65,535
Default: 0

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Falling Event Index or Rising Event Index on page 243](#)
- [rising-event-index on page 701](#)

falling-threshold

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	falling-threshold <i>integer</i> ;
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the lower threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>integer</i> —The lower threshold for the alarm entry. Range: -2,147,483,648 through 2,147,483,647 Default: 20 percent less than rising-threshold
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Falling Threshold or Rising Threshold on page 244• rising-threshold on page 702

falling-threshold-interval

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `falling-threshold-interval seconds;`

Hierarchy Level [edit snmp rmon [alarm index](#)]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Set the interval between samples after the rising threshold is exceeded and the value of the sample starts to drop. If the value of the sample drops and exceeds the falling threshold, the regular sampling interval is used.

Options *seconds*—Time between samples, in seconds.
Range: 1 through 2,147,483,647 seconds
Default: 60 seconds

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Falling Threshold Interval on page 244](#)
- [interval on page 699](#)

interval

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	interval <i>seconds</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Interval between samples.
Options	seconds —Time between samples, in seconds. Range: 1 through 2,147,483,647 seconds Default: 60 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interval on page 244

request-type

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax request-type (get-next-request | get-request | walk-request);

Hierarchy Level [edit snmp rmon [alarm index](#)]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Extend monitoring to a specific SNMP object instance (**get-request**), or extend monitoring to all object instances belonging to a MIB branch (**walk-request**), or extend monitoring to the next object instance after the instance specified in the configuration (**get-next-request**).

Options **get-next-request**—Performs an SNMP get next request.

get-request—Performs an SNMP get request.

walk-request—Performs an SNMP walk request.

Default: walk-request

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Request Type on page 245](#)
- [variable on page 707](#)

rising-event-index

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	rising-event-index <i>index</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Index of the event entry that is used when a rising threshold is crossed. If this value is zero, no event is triggered.
Options	<i>index</i> —Index of the event entry that is used when a rising threshold is crossed. Range: 0 through 65,535 Default: 0
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Falling Event Index or Rising Event Index on page 243• falling-event-index on page 696

rising-threshold

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	rising-threshold <i>integer</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the upper threshold for the sampled variable (monitored object) when you configure an SNMP RMON alarm. By setting a rising and a falling threshold for a variable, you can be alerted whenever the value of the variable falls outside the allowable operational range.
Options	<i>integer</i> —The lower threshold for the alarm entry. Range: –2,147,483,648 through 2,147,483,647
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Falling Threshold or Rising Threshold on page 244• falling-threshold on page 697

rmon

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	rmon { ... }
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure Remote Monitoring.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an RMON Alarm Entry and Its Attributes on page 242

sample-type

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	sample-type (absolute-value delta-value);
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Method of sampling the selected variable.
Options	<p>absolute-value—Actual value of the selected variable is used when comparing against the thresholds.</p> <p>delta-value—Difference between samples of the selected variable is used when comparing against the thresholds.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Sample Type on page 245

startup-alarm

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	startup-alarm (falling-alarm rising-alarm rising-or-falling-alarm);
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The alarm that can be sent upon entry startup.
Options	<p>falling-alarm—Generated if the first sample after the alarm entry becomes active is less than or equal to the falling threshold.</p> <p>rising-alarm—Generated if the first sample after the alarm entry becomes active is greater than or equal to the rising threshold.</p> <p>rising-or-falling-alarm—Generated if the first sample after the alarm entry becomes active satisfies either of the corresponding thresholds.</p> <p>Default: rising-or-falling-alarm</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Startup Alarm on page 245

syslog-subtag

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	syslog-subtag <i>syslog-subtag</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Add a tag to the system log message.
Options	syslog-subtag <i>syslog-subtag</i> —Tag of not more than 80 uppercase characters to be added to syslog messages. Default: None
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the System Log Tag on page 246

type

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	type type;
Hierarchy Level	[edit snmp rmon event index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Type of notification generated when a threshold is crossed.
Options	<p>type—Type of notification:</p> <ul style="list-style-type: none">• log—Add an entry to logTable.• log-and-trap—Send an SNMP trap and make a log entry.• none—No notifications are sent.• snmptrap—Send an SNMP trap. <p>Default: log-and-trap</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an RMON Event Entry and Its Attributes on page 246

variable

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	<code>variable <i>oid-variable</i>;</code>
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Object identifier (OID) of MIB variable to be monitored.
Options	<i>oid-variable</i> —OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1). Alternatively, use the MIB object name (for example, <code>ifInOctets.1</code>).
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Variable on page 246

CHAPTER 40

Configuration Statements: Resource Monitoring for Memory Regions

- `free-fw-memory-watermark` (Resource Monitor) on page 710
- `free-heap-memory-watermark` (Resource Monitor) on page 711
- `free-nh-memory-watermark` (Resource Monitor) on page 712
- `high-cos-queue-threshold` on page 713
- `high-threshold` (Resource Monitor) on page 714
- `no-load-throttle` (Resource Monitor) on page 714
- `no-logging` (Resource Monitor) on page 715
- `no-throttle` (Resource Monitor) on page 716
- `resource-category jtree` (Resource Monitor) on page 717
- `resource-monitor` on page 718
- `subscribers-limit` (Resource Monitor) on page 720
- `traceoptions` (Resource Monitor) on page 721

free-fw-memory-watermark (Resource Monitor)

Supported Platforms [MX104](#), [MX2010](#), [MX2020](#), [MX240](#), [MX480](#), [MX80](#), [MX960](#)

Syntax `free-fw-memory-watermark number;`

Hierarchy Level [edit system services resource-monitor]

Release Information Statement introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Configure the percentage of free memory space used for firewall or filters to be monitored with a watermark value. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options *number*—Percentage of free memory space used for firewall and filters to be monitored with a watermark value. When the configured watermark is exceeded, error logs are triggered. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20.
Range: 1 through 100

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

free-heap-memory-watermark (Resource Monitor)

Supported Platforms [MX104](#), [MX2010](#), [MX2020](#), [MX240](#), [MX480](#), [MX80](#), [MX960](#)

Syntax `free-heap-memory-watermark number;`

Hierarchy Level [edit system services resource-monitor]

Release Information Statement introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Configure the percentage of free memory space used for ukernel or heap (ASIC) memory to be monitored with a watermark value. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options *number*—Percentage of free memory space used for ukernel or heap to be monitored with a watermark value. When the configured watermark is exceeded, error logs are triggered. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20.
Range: 1 through 100

Required Privilege system—To view this statement in the configuration.
Level system-control—To add this statement to the configuration.

free-nh-memory-watermark (Resource Monitor)

Supported Platforms	MX104, MX2010, MX2020, MX240, MX480, MX80, MX960
Syntax	free-nh-memory-watermark <i>number</i> ;
Hierarchy Level	[edit system services resource-monitor]
Release Information	Statement introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.
Description	Configure the percentage of free memory space used for next-hops to be monitored with a watermark value. The default value and the configured value of the watermark value for the percentage of free next-hop memory also applies to encapsulation memory. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.
Options	<p><i>number</i>—Percentage of free memory space used for next-hops to be monitored with a watermark value. The NH memory watermark is applicable only for encapsulation memory (output WAN static RAM memory). When the configured watermark is exceeded, error logs are triggered. The default watermark values for the percentage of free ukernel or heap memory, next-hop memory, and firewall filter memory are 20.</p> <p>Range: 1 through 100</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>

high-cos-queue-threshold

Supported Platforms MX2010, MX2020, MX240, MX480, MX960

Syntax high-cos-queue-threshold *number*;

Hierarchy Level [edit system services [resource-monitor](#)]

Release Information Statement introduced in Junos OS Release 17.4R1.

Description Configure the class-of-service (CoS) queue threshold by percentage.



NOTE: This feature is only available when you enable subscriber management. For more information on enabling subscriber management, see *Configuring Junos OS Enhanced Subscriber Management*.

This CoS resource monitoring feature bases admission decisions only on queues. Other CoS resources are not part of this criteria. This feature does not support throttling for subscribers arriving on pseudo-wire, logical tunnel, or redundant logical tunnel devices. The feature is supported on the following hardware:

- MX240, MX480, and MX960 routers
- MPC2E legacy, MPC2E-NG, MPC3E-NG, MPC5E, and MPC7E line cards

Options **number**—High-threshold percentage for CoS queue utilization per scheduler. When CoS queue utilization on a given FPC reaches that FPC's configured threshold level, further subscriber logins on that FPC are not allowed. This resource monitoring mechanism provides adjustable safety margins to proactively avoid completely exhausting each FPC's available CoS queue resources.

Range: 1 through 90

Default: 100

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *Resource Monitoring for Subscriber Management and Services Overview*

high-threshold (Resource Monitor)

Supported Platforms	MX104, MX2010, MX2020, MX240, MX480, MX80, MX960
Syntax	high-threshold <i>number</i> ;
Hierarchy Level	[edit system services resource-monitor]
Release Information	Statement introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.
Description	Configure the high threshold value, exceeding which warnings or error logs are generated, for all the regions of memory, such as heap or ukernel, next-hop and encapsulation, and firewall filter memory. You can configure the resource-monitoring capability on MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.
Options	<i>number</i> —High threshold percentage for memory resource utilization Range: 1 through 100
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

no-load-throttle (Resource Monitor)

Supported Platforms	MX Series
Syntax	no-load-throttle;
Hierarchy Level	[edit system services resource-monitor]
Release Information	Statement introduced in Junos OS Release 17.3R1 for MX series routers.
Description	The no-load-throttle statement disables line card load-based throttling. Load-based throttling is also disabled when you configure the no-throttle statement.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

no-logging (Resource Monitor)

Supported Platforms	MX104, MX2010, MX2020, MX240, MX480, MX80, MX960
Syntax	no-logging;
Hierarchy Level	[edit system services resource-monitor]
Release Information	Statement introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.
Description	Disable the generation of error log messages when the utilization of memory resources exceeds the threshold or checkpoint levels. By default, messages are written to /var/log/rsmonlog.
Options	no-logging —Disable the generation of error log messages when the utilization of memory resources exceeds the configured level. By default, error logs are recorded when the resource level utilization is exceeded.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

no-throttle (Resource Monitor)

Supported Platforms [MX2010](#), [MX2020](#), [MX240](#), [MX480](#), [MX960](#)

Syntax no-throttle;

Hierarchy Level [edit system services resource-monitor]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Disable the throttling of subscriber services and sessions when the utilization of memory resources exceeds the threshold levels. For the subscriber service management, sessions are throttled based on line card resources, such as heap, counter memory, expansion memory, and firewall counter memory. If the resource has been used above a certain threshold, the subscribers and services are throttled to prevent the system from being overloaded and resulting in a breakdown. This feature gathers input from each of the line cards and transfers this statistical detail to the Routing Engine process using a well-known internal port. This information is scanned by the daemon on the Routing Engine and using the shared memory space built into the session database, the existing active subscribers and sessions are throttled. You can configure the resource-monitoring capability on MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

Options **no-throttle**—Disable the throttling of subscriber services and sessions when the utilization of memory resources exceeds the threshold levels. The throttling capability is enabled by default.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

resource-category jtree (Resource Monitor)

Supported Platforms MX2010, MX2020, MX240, MX480, MX960

Syntax resource-category jtree {
 resource-type (contiguous-pages | free-dwords | free-pages) {
 low-watermark *number*;
 high-watermark *number*;
 }
 }

Hierarchy Level [edit system services resource-monitor]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Define the resource category that you want to monitor and analyze for ensuring system stability, especially the health and operating efficiency of I-chip-based line cards and Trio-based FPCs on MX Series routers. The resource category includes detailed CPU utilization, session rate, and session count statistics. You use the resource category statistics to understand the extent to which new attack objects or applications affect performance. The jtree memory on all MX Series router Packet Forwarding Engines has two segments: one segment primarily stores routing tables and related information, and the other segment primarily stores firewall-filter-related information. The Junos OS provides the memory-enhanced statement to reallocate the jtree memory for routes, firewall filters, and Layer 3 VPNs.

Options **jtree**—Specify the Jtree resource category for which you want to monitor the health (working condition), operating efficiency, traffic-handling capacity, and performance status.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level system—To view this statement in the configuration.
 system-control—To add this statement to the configuration.

resource-monitor

Supported Platforms MX2010, MX2020, MX240, MX480, MX960

Syntax

```
resource-monitor {
  free-fw-memory-watermark number;
  free-heap-memory-watermark number;
  free-nh-memory-watermark number;
  high-cos-queue-threshold number;
  high-threshold number;
  no-logging;
  no-throttle;
  resource-category jtree {
    resource-type (contiguous-pages | free-dwords | free-pages) {
      low-watermark number;
      high-watermark number;
    }
  }
  subscribers-limit {
    client-type (any | dhcp | l2tp | pppoe) {
      chassis {
        limit limit;
      }
      fpc slot-number {
        limit limit;
        pic number {
          limit limit;
          port number {
            limit limit;
          }
        }
      }
    }
  }
  traceoptions {
    file filename <files number> <match regular-expression> <size maximum-file-size>
      <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced in Junos OS Release 15.1 for MX240, MX480, MX960, MX2010, and MX2020 routers.
high-cos-queue-threshold option introduced in Junos OS Release 17.4R1.

Description Enable the resource monitoring capability to provision sufficient headroom (memory space limits that are set for the application or virtual router) for monitoring the health and operating efficiency of DPCs and MPCs. This feature also enables the memory resource monitoring mechanism to avoid the system operations from compromising on

the health and traffic-handling stability of the line cards by generating error logs when a specified watermark value for memory regions and threshold value for the jtree memory region are exceeded. A trade-off on the system performance can be detrimental for supporting live traffic and protocols.

The variable *number* in the Syntax section represents a percentage.

You can configure the resource-monitoring capability on MX240, MX480, MX960, MX2010, and MX2020 routers with I-chip-based DPCs and Trio-based FPCs.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Diagnosing and Debugging System Performance by Configuring Memory Resource Usage Monitoring on MX Series Routers on page 520• <i>Resource Monitoring for Subscriber Management and Services Overview</i>• Resource Monitoring Usage Computation Overview on page 517• <i>Limiting Subscribers by Client Type and Hardware Element with Resource Monitor</i>

subscribers-limit (Resource Monitor)

Supported Platforms [MX Series](#)

Syntax

```
subscribers-limit {
  client-type (any | dhcp | l2tp | pppoe) {
    chassis {
      limit limit;
    }
    fpc slot-number {
      limit limit;
      pic number {
        limit limit;
        port number {
          limit limit;
        }
      }
    }
  }
}
```

Hierarchy Level [edit system services [resource-monitor](#)]

Release Information Statement introduced in Junos OS Release 17.3R1.

Description Configure the maximum number of subscribers of a specified client type allowed to be logged in on the chassis, per MPC, per MIC, and per port. When that number is reached, subsequent logins are denied until the current number of subscribers drops below the maximum allowed.

Limit the number of subscribers allowed to log in per chassis, MPC, MIC, or port.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system	—To view this statement in the configuration.
system-control	—To add this statement to the configuration.

Related Documentation

- *Limiting Subscribers by Client Type and Hardware Element with Resource Monitor*
- *Resource Monitoring for Subscriber Management and Services Overview*

tracoptions (Resource Monitor)

Supported Platforms MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

Syntax tracoptions {
 file *filename* <files *number*> <match *regular-expression* > <size *maximum-file-size*>
 <world-readable | no-world-readable>;
 flag *flag*;
 }

Hierarchy Level [edit system services resource-monitor]

Release Information Statement introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.

Description Define tracing operations for the memory resource utilization processes.

Options file *filename*—Name of the file to receive the output of the tracing operation. All files are placed in the directory `/var/log`.

Default: rmopd

files *number*—(Optional) Maximum number of trace files to create before overwriting the oldest one. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

size *maximum-file-size*—(Optional) Maximum size of each trace file. By default, the number entered is treated as bytes. Alternatively, you can include a suffix to the number to indicate kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

no-world-readable—(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Trace all operations.

Required Privilege trace—To view this statement in the configuration.
Level trace-control—To add this statement to the configuration.

Configuration Statements: Security Alarms

- [decryption-failures](#) on page 723
- [idp \(Security Alarms\)](#) on page 724

decryption-failures

Supported Platforms [SRX Series, vSRX](#)

Syntax `decryption-failures {
 threshold value;
}`

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Raise a security alarm after exceeding a specified number of decryption failures. This statement is supported on SRX300, SRX320, SRX340, SRX345, SRX550HM, and SRX1500 devices and vSRX instances.

Default Multiple decryption failures do not cause an alarm to be raised.

Options *failures*—Number of decryption failures up to which an alarm is not raised. When the configured number is exceeded, an alarm is raised.

Range: 0 through 1 through 1,000,000,000.

Default: 1000

Required Privilege Level security—To view this statement in the configuration.
security-control—To add this statement to the configuration.

Related Documentation

- [IPsec VPN Overview](#)

idp (Security Alarms)

Supported Platforms [SRX Series, vSRX](#)

Syntax `idp;`

Hierarchy Level [edit security alarms potential-violation]

Release Information Statement introduced in Junos OS Release 11.2.

Description Configure alarms for IDP attack.

Required Privilege security—To view this statement in the configuration.
Level security-control—To add this statement to the configuration.

Configuration Statements: SNMP

- [access-list on page 726](#)
- [agent-address on page 727](#)
- [alarm-id on page 728](#)
- [alarm-list-name on page 729](#)
- [alarm-management on page 730](#)
- [alarm-state on page 731](#)
- [authorization on page 732](#)
- [categories on page 733](#)
- [client-list on page 734](#)
- [client-list-name on page 735](#)
- [clients on page 736](#)
- [commit-delay on page 737](#)
- [community \(SNMP\) on page 738](#)
- [contact \(SNMP\) on page 739](#)
- [description on page 740](#)
- [destination-port on page 741](#)
- [enterprise-oid on page 741](#)
- [filter-duplicates on page 742](#)
- [filter-interfaces on page 743](#)
- [interface \(SNMP\) on page 744](#)
- [location \(SNMP\) on page 745](#)
- [logical-system on page 746](#)
- [logical-system-trap-filter on page 747](#)
- [name on page 747](#)
- [nonvolatile on page 748](#)
- [oid on page 749](#)
- [proxy \(snmp\) on page 750](#)
- [routing-instance on page 752](#)

- [routing-instance-access](#) on page 753
- [snmp](#) on page 753
- [snmp-value-match-msmic](#) (Services NAT Options) on page 754
- [source-address](#) on page 755
- [targets](#) on page 756
- [traceoptions](#) (SNMP) on page 757
- [trap-group](#) on page 759
- [trap-options](#) on page 760
- [version](#) (SNMP) on page 761
- [view](#) (Associating a MIB View with a Community) on page 762
- [view](#) (Configuring a MIB View) on page 763

access-list

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax

```
[edit snmp]
  routing-instance-access {
    access-list {
      routing-instance;
      routing-instance restrict;
    }
  }
```

Hierarchy Level `[edit snmp routing-instance-access]`

Release Information Statement introduced in Junos OS Release 8.4.

Description Create access lists to control SNMP agents in routing instances from accessing SNMP information. To enable the SNMP agent on a routing instance to access SNMP information, specify the routing instance name. To disable the SNMP agent on a routing instance from accessing SNMP information, include the routing-instance name followed by the **restrict** keyword.

Required Privilege `snmp`—To view this statement in the configuration.
Level `snmp-control`—To add this statement to the configuration.

Related Documentation

- [routing-instance-access](#) on page 753

agent-address

Supported Platforms	ACX Series , EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	agent-address outgoing-interface;
Hierarchy Level	[edit snmp trap-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is outgoing-interface , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
Options	outgoing-interface —Value of the agent address of all SNMPv1 traps generated by this router or switch. The outgoing-interface option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. Default: disabled (the agent address is not specified in SNMPv1 traps).
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Agent Address for SNMP Traps on page 105

alarm-id

Supported Platforms [MX Series](#)

Syntax `alarm-id id {
 alarm-state state {
 description alarm-description;
 notification-id notification-id-of-alarm;
 resource-prefix alarm-resource-prefix;
 varbind-index varbind-index-in-alarm-varbind-list;
 varbind-subtree alarm-varbind-subtree;
 varbind-value alarm-varbind-value;
 }
 }`

Hierarchy Level `[edit snmp alarm-management alarm-list-name]`

Release Information Statement introduced in Junos OS Release 14.1.

Description Specify the identifier of the alarm that you need to configure.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level `snmp`—To view this statement in the configuration.
 `snmp-control`—To add this statement to the configuration.

Related Documentation

- [alarm-list-name on page 729](#)
- [alarm-management on page 730](#)
- [alarm-state on page 731](#)

alarm-list-name

Supported Platforms [MX Series](#)

Syntax

```
alarm-list-name list-name {
  alarm-id id {
    alarm-state state {
      description alarm-description;
      notification-id notification-id-of-alarm;
      resource-prefix alarm-resource-prefix;
      varbind-index varbind-index-in-alarm-varbind-list;
      varbind-subtree alarm-varbind-subtree;
      varbind-value alarm-varbind-value;
    }
  }
}
```

Hierarchy Level [edit snmp alarm-management]

Release Information Statement introduced in Junos OS Release 14.1.

Description Specify the name of the alarm list that you need to configure.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [alarm-id on page 728](#)
- [alarm-management on page 730](#)
- [alarm-state on page 731](#)

alarm-management

Supported Platforms [MX Series](#)

Syntax

```
alarm-management {
  alarm-list-name list-name {
    alarm-id id {
      alarm-state state {
        description alarm-description;
        notification-id notification-id-of-alarm;
        resource-prefix alarm-resource-prefix;
        varbind-index varbind-index-in-alarm-varbind-list;
        varbind-subtree alarm-varbind-subtree;
        varbind-value alarm-varbind-value;
      }
    }
  }
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure the alarm management system to monitor and report active alarms as well as the history of alarms through the SNMP MIB tables supported by the *Alarm MIB*.



NOTE: You cannot configure alarms without notifications. It is mandatory to include the notification identifier in the configuration.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [alarm-id on page 728](#)
- [alarm-list-name on page 729](#)
- [alarm-state on page 731](#)

alarm-state

Supported Platforms [MX Series](#)

Syntax `alarm-state state {
 description alarm-description;
 notification-id notification-id-of-alarm;
 resource-prefix alarm-resource-prefix;
 varbind-index varbind-index-in-alarm-varbind-list;
 varbind-subtree alarm-varbind-subtree;
 varbind-value alarm-varbind-value;
}`

Hierarchy Level [edit snmp alarm-management alarm-list-name]

Release Information Statement introduced in Junos OS Release 14.1.

Description Specify the state of the alarm and the other parameters that you need to monitor.

Options

- description** *alarm-description*—Include a brief description of the alarm.
- notification-id** *notification-id-of-alarm*—Specify the identifier of the notification associated with the alarm.
- resource-prefix** *alarm-resource-prefix*—Specify the resource prefix of the alarm.
- varbind-index** *varbind-index-in-alarm-varbind-list*—Specify the varbind index in the alarm varbind list.
 Range: 0 through 4294967295
- varbind-subtree** *alarm-varbind-subtree*—Specify the subtree of the varbind.
- varbind-value** *alarm-varbind-value*—Specify the varbind value of the alarm.
 Range: 0 through 2147483647

Required Privilege Level

- snmp—To view this statement in the configuration.
- snmp-control—To add this statement to the configuration.

Related Documentation

- [alarm-id on page 728](#)
- [alarm-list-name on page 729](#)
- [alarm-management on page 730](#)

authorization

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric System, QFX Series, SRX Series, T Series
Syntax	authorization <i>authorization</i> ;
Hierarchy Level	[edit snmp community <i>community-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the access authorization for SNMP Get , GetBulk , GetNext , and Set requests.
Options	<i>authorization</i> —Access authorization level: <ul style="list-style-type: none">• read-only—Enable Get, GetNext, and GetBulk requests.• read-write—Enable all requests, including Set requests. You must configure a view to enable Set requests. Default: read-only
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Communities on page 93

categories

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric System, QFX Series, SRX Series, T Series
Syntax	<pre>categories { category; }</pre>
Hierarchy Level	[edit snmp trap-group <i>group-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Define the types of traps that are sent to the targets of the named trap group.
Default	If you omit the categories statement, all trap types are included in trap notifications.
Options	category —Name of a trap type: authentication , chassis , configuration , link , remote-operations , rmon-alarm , routing , services , sonet-alarms , startup , or vrrp-events .
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Trap Groups on page 106

client-list

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	<pre>client-list <i>client-list-name</i> { <i>ip-addresses</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series switches.
Description	Define a list of SNMP clients.
Options	<i>client-list-name</i> —Name of the client list. <i>ip-addresses</i> —IP addresses of the SNMP clients to be added to the client list,
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Adding a Group of Clients to an SNMP Community on page 97

client-list-name

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `client-list-name client-list-name;`

Hierarchy Level [edit snmp community *community-name*]

Release Information Statement introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Add a client list or prefix list to an SNMP community.

Options *client-list-name*—Name of the client list or prefix list.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Adding a Group of Clients to an SNMP Community on page 97](#)

clients

Supported Platforms	ACX Series , EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	<pre>clients { address <restrict>; }</pre>
Hierarchy Level	[edit snmp community <i>community-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
Default	If you omit the clients statement, all SNMP clients using this community string are authorized to access the router.
Options	address —Address of an SNMP client that is authorized to access this router. You must specify an address, not a hostname. To specify more than one client, include multiple address options. restrict —(Optional) Do not allow the specified SNMP client to access the router.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Communities on page 93

commit-delay

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric System, QFX Series, SRX Series, T Series
Syntax	commit-delay <i>seconds</i> ;
Hierarchy Level	[edit snmp nonvolatile]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the timer for the SNMP Set reply and start of the commit.
Options	seconds —Delay between an affirmative SNMP Set reply and start of the commit. Default: 5 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Commit Delay Timer on page 92

community (SNMP)

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `community community-name {
 authorization authorization;
 client-list-name client-list-name;
 clients {
 address restrict;
 }
 view view-name;
}`

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.

The SNMP client application specifies an SNMP community name in **Get**, **GetBulk**, **GetNext**, and **Set** SNMP requests.

Default If you omit the **community** statement, all SNMP requests are denied.

Options ***community-name***—Community string. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Communities on page 93](#)

contact (SNMP)

Supported Platforms	ACX Series , EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	<code>contact <i>contact</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define the value of the MIB II sysContact object, which is the contact person for the managed system.
Options	contact —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the System Contact on a Device Running Junos OS on page 88

description

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series, vSRX
Syntax	<code>description <i>description</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysDescription object, which is the description of the system being managed.
Options	<i>description</i> —System description. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the System Description on a Device Running Junos OS on page 89

destination-port

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `destination-port port-number;`

Hierarchy Level [edit snmp trap-group]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Assign a trap port number other than the default.

Default If you omit this statement, the default port is 162.

Options *port-number*—SNMP trap port number.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Trap Groups on page 106](#)

enterprise-oid

Supported Platforms [ACX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `enterprise-oid;`

Hierarchy Level [edit snmp [trap-options](#)]

Release Information Statement introduced in Junos OS Release 10.0

Description Add the `snmpTrapEnterprise` object, which shows the association between an enterprise-specific trap and the organization that defined the trap, to standard SNMP traps. By default, the `snmpTrapEnterprise` object is added only to the enterprise-specific traps. When the `enterprise-oid` statement is included in the configuration, `snmpTrapEnterprise` is added to all the traps generated from the device.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Trap Options on page 103](#)

filter-duplicates

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax filter-duplicates;

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Filter duplicate **Get**, **GetNext**, or **GetBulk** SNMP requests.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Filtering Duplicate SNMP Requests on page 93](#)

filter-interfaces

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax

```
filter-interfaces {
  interfaces {
    all-internal-interfaces;
    interface 1;
    interface 2;
  }
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series Switches.

Description Filter out information related to specific interfaces from the output of SNMP **Get** and **GetNext** requests performed on interface-related MIBs.

Options

- all-internal-interfaces**—Filters out information from SNMP **Get** and **GetNext** requests for the specified interfaces.
- interfaces**—Specifies the interfaces to filter out from the output of SNMP **Get** and **GetNext** requests.

Required Privilege Level

- snmp—To view this statement in the configuration.
- snmp-control—To add this statement to the configuration.

Related Documentation

- [Filtering Interface Information Out of SNMP Get and GetNext Output on page 110](#)

interface (SNMP)

Supported Platforms	ACX Series, EX Series, M120, MX240, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T1600, T640
Syntax	interface [<i>interface-names</i>];
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Configure the interfaces on which SNMP requests can be accepted.
Default	If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.
Options	<i>interface-names</i> —Names of one or more logical interfaces.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interfaces on Which SNMP Requests Can Be Accepted on page 109

location (SNMP)

Supported Platforms	ACX Series , EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	location <i>location</i> ;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define the value of the MIB II sysLocation object, which is the physical location of the managed system.
Options	location —Location of the local system. You must enclose the name within quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the System Location for a Device Running Junos OS on page 89

logical-system

Supported Platforms EX Series, M120, MX240, PTX Series, SRX Series, SRX320, T1600

Syntax `logical-system logical-system-name {
 routing-instance routing-instance-name;
 source-address address;
}`

Hierarchy Level [edit snmp *community community-name*],
[edit snmp *trap-group*],
[edit snmp *trap-options*]
[edit snmp v3 *target-address target-address-name*]

Release Information Statement introduced in Junos OS Release 9.3
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.



NOTE: The `logical-system` statement replaces the `logical-router` statement, and is backward-compatible with Junos OS Release 8.3 and later.

Description Specify a logical system name for SNMP v1 and v2c clients.

Include at the [edit snmp *trap-options*] hierarchy level to specify a logical-system address as the source address of an SNMP trap.

Include at the [edit snmp v3 *target-address*] hierarchy level to specify a logical-system name as the destination address for an SNMPv3 trap or inform.

Options *logical-system-name*—Name of the logical system.

routing-instance routing-instance-name—Statement to specify a routing instance associated with the logical system.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169](#)
- [Configuring the Trap Target Address on page 137](#)

logical-system-trap-filter

Supported Platforms	M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	logical-system-trap-filter;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 8.4.
Description	Restrict the routing instances from receiving traps that are not related to the logical system networks to which they belong.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • SNMP Traps Supported for Routing Instances on page 167

name

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	name <i>name</i> ;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the system name from the command-line interface.
Options	<i>name</i> —System name override.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring a Different System Name on page 91

nonvolatile

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	<pre>nonvolatile { commit-delay seconds; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. The commit-delay statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure options for SNMP Set requests. The remaining statement is explained separately. See CLI Explorer .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Commit Delay Timer on page 92• commit-delay on page 737

oid

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#), [vSRX](#)

Syntax `oid object-identifier (exclude | include);`

Hierarchy Level `[edit snmp view view-name]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify an object identifier (OID) used to represent a subtree of MIB objects.

Options **exclude**—Exclude the subtree of MIB objects represented by the specified OID.

include—Include the subtree of MIB objects represented by the specified OID.

object-identifier—OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.

Required Privilege Level `snmp`—To view this statement in the configuration.
`snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring MIB Views on page 111](#)

proxy (snmp)

Supported Platforms [M Series](#), [MX Series](#), [T Series](#)

Syntax

```
proxy proxy-name{
  device-name device-name;
  logical-system logical-system {
    routing-instance routing-instance;
  }
  routing-instance routing-instance;
  (version-v1 | version-v2c) {
    snmp-community community-name;
    no-default-comm-to-v3-config;
  }
  version-v3 {
    security-name security-name;
    context context-name;
  }
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 12.3.

Description Configure a device to act as a proxy SNMP agent, and specify a name for the proxy.

Options **context** *context-name*—Specify the SNMPv3 context name as configured on the device specified at **edit snmp proxy** *proxy-name* **device-name** *device-name*.

For more information about this statement, see [context](#).

device-name *device-name*—Specify the name of the device to be managed through the proxy SNMP agent.

no-default-comm-to-v3-config—(Optional) Specify whether you have to manually configure the statements at the **[edit snmp v3 snmp-community *community-name*]** and **[edit snmp v3 vacm]** hierarchy levels.

If this statement is not included in the configuration, the **[edit snmp v3 snmp-community *community-name*]** and **[edit snmp v3 vacm]** hierarchy level configurations are automatically initialized.

proxy-name—Specify the name of the proxy.

security-name *security-name*—Specify the SNMPv3 security name as configured on the device specified at **edit snmp proxy** *proxy-name* **device-name** *device-name*.

For more information about this statement, see [security-name](#).

snmp-community *community-name*—Specify the name of the SNMP community. The community name you configure should match the **snmp-community** configuration on the device specified at **edit snmp proxy *proxy-name* device-name *device-name***. For more information about this statement, see [snmp-community](#).

(version-v1 | version-v2c)—Specify the SNMP version, and add the relevant configuration.

version-v3—Add the SNMPv3 configuration.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring a Proxy SNMP Agent on page 98
------------------------------	---

routing-instance

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<code>routing-instance <i>routing-instance-name</i>;</code>
Hierarchy Level	[edit snmp community <i>community-name</i>], [edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>], [edit snmp trap-group <i>group</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Added to the [edit snmp community <i>community-name</i>] hierarchy level in Junos OS Release 8.4. Added to the [edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>] hierarchy level in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
Description	<p>Specify a routing instance for SNMPv1 and SNMPv2 trap targets. All targets configured in the trap group use this routing instance.</p> <p>If the routing instance is defined within a logical system, include the logical-system <i>logical-system-name</i> statement at the [edit snmp community <i>community-name</i>] hierarchy level and specify the routing-instance statement under the [edit snmp community <i>community-name</i> logical-system <i>logical system-name</i>] hierarchy level.</p>
Options	<i>routing-instance-name</i> —Name of the routing instance.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Trap Groups on page 106• Configuring the Source Address for SNMP Traps on page 103• Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community on page 169

routing-instance-access

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax

```
[edit snmp]
  routing-instance-access {
    access-list {
      routing-instance;
      routing-instance restrict;
    }
  }
```

Hierarchy Level [edit snmp]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable SNMP managers in routing instances other than the default routing instance to access SNMP information. For information about the **access-list** option, see [access-list](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Enabling SNMP Access over Routing Instances on page 169](#)

snmp

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Syntax snmp { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure SNMP.

SNMP modules cannot have the slash (/) character or the @ character in the name.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP on a Device Running Junos OS](#)

snmp-value-match-msmic (Services NAT Options)

Supported Platforms [MX Series](#)

Syntax `snmp-value-match-msmic;`

Hierarchy Level `[edit services service-set service-set-name nat-options]`

Release Information Statement introduced with Junos OS Release 13.3R7

Description Match the values for MS-MIC-specific objects in the jnxNatObjects MIB table with the values for MS-DPC objects. Use the **deactivate services service-set *service-set-name* nat-options snmp-value-match-msmic** configuration mode command to disable this feature. By default, this feature is disabled.

Required Privilege Level control—To add this statement to the configuration.

Related Documentation

- [Configuring Service Rules](#)
- [Troubleshooting the Mismatch of jnxNatObjects Values for MS-DPC and MS-MIC on page 523](#)

source-address

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	source-address <i>address</i> ;
Hierarchy Level	[edit snmp trap-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the source address of every SNMP trap packet sent by this router to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.
Options	address —Source address of SNMP traps. You can configure the source address of trap packets two ways: lo0 or a valid IPv4 address configured on one of the router interfaces. The value lo0 indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface lo0 . Default: Disabled. (The source address is the address of the outgoing interface.)
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Source Address for SNMP Traps on page 103

targets

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax

```
targets {  
    address;  
}
```

Hierarchy Level [edit snmp trap-group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure one or more systems to receive SNMP traps.

Options *address*—IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Trap Groups on page 106](#)

traceoptions (SNMP)

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX320, T Series, vSRX
Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; no-remote-trace; } </pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>file <i>filename</i> option added in Junos OS Release 8.1.</p> <p>world-readable no-world-readable option added in Junos OS Release 8.1.</p> <p>match <i>regular-expression</i> option added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>The output of the tracing operations is placed into log files in the /var/log directory. Each log file is named after the SNMP agent that generates it. Currently, the following logs are created in the /var/log directory when the traceoptions statement is used:</p> <ul style="list-style-type: none"> • chassisd • craftd • ilmids • mib2d • rmopd • serviced • snmpd
Options	<p>file <i>filename</i>—By default, the name of the log file that records trace output is the name of the process being traced (for example, mib2d or snmpd). Use this option to specify another name.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, snmpd) reaches its maximum size, it is archived by being renamed to snmpd.0. The previous snmpd.1 is renamed to snmpd.2, and so on. The oldest archived file is deleted.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements:</p>

- **all**—Log all SNMP events.
- **general**—Log general events.
- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **subagent**—Log subagent restarts.
- **timer**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

size *size*—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

Range: 10 KB through 1 GB

Default: 1000 KB

world-readable | no-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Tracing SNMP Activity on a Device Running Junos OS on page 201
------------------------------	--

trap-group

Supported Platforms	ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series, vSRX
Syntax	<pre> trap-group group-name { categories { category; } destination-port port-number; routing-instance instance; targets { address; } version (all v1 v2); } </pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for OCX Series switches.</p>
Description	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
Options	<p>group-name—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP Trap Groups on page 106

trap-options

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax

```
trap-options {  
    agent-address outgoing-interface;  
    source-address address;  
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.

The remaining statements are explained separately. See [CLI Explorer](#).

Default Disabled

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Trap Options on page 103](#)

version (SNMP)

Supported Platforms [ACX Series](#), [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `version (all | v1 | v2);`

Hierarchy Level `[edit snmp trap-group group-name]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify the version number of SNMP traps.

Default `all`—Send an SNMPv1 and SNMPv2 trap for every trap condition.

Options `all`—Send an SNMPv1 and SNMPv2 trap for every trap condition.

`v1`—Send SNMPv1 traps only.

`v2`—Send SNMPv2 traps only.

Required Privilege Level `snmp`—To view this statement in the configuration.
`snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Trap Groups on page 106](#)

view (Associating a MIB View with a Community)

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	<code>view view-name;</code>
Hierarchy Level	[edit snmp community <i>community-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Associate a view with a community. A view represents a group of MIB objects.
Options	<i>view-name</i> —Name of the view. You must use a view name already configured in the view statement at the [edit snmp] hierarchy level.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Communities on page 93

view (Configuring a MIB View)

Supported Platforms ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series

Syntax `view view-name {
oid object-identifier (include | exclude);
}`

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The **view** statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the **view** statement at the [edit snmp community *community-name*] hierarchy level.



NOTE: To remove an OID completely, use the `delete view all oid oid-number` command but omit the `include` parameter.

Options *view-name*—Name of the view.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring MIB Views on page 111](#)
- [Associating MIB Views with an SNMP User Group on page 129](#)
- [community on page 738](#)

Configuration Statements: SNMPv3

- [address on page 767](#)
- [address-mask on page 767](#)
- [authentication-key on page 768](#)
- [authentication-md5 on page 769](#)
- [authentication-none on page 770](#)
- [authentication-password on page 771](#)
- [authentication-sha on page 772](#)
- [community-name on page 773](#)
- [context \(SNMPv3\) on page 774](#)
- [engine-id on page 775](#)
- [group \(Configuring Group Name\) on page 776](#)
- [group \(Defining Access Privileges for an SNMPv3 Group\) on page 777](#)
- [retry-count on page 778](#)
- [timeout on page 779](#)
- [local-engine on page 780](#)
- [message-processing-model on page 781](#)
- [notify on page 782](#)
- [notify-filter \(Applying to the Management Target\) on page 783](#)
- [notify-filter \(Configuring the Profile Name\) on page 784](#)
- [notify-view on page 785](#)
- [oid on page 786](#)
- [parameters on page 787](#)
- [port on page 788](#)
- [privacy-3des on page 789](#)
- [privacy-aes128 on page 790](#)
- [privacy-des on page 791](#)
- [privacy-key on page 792](#)
- [privacy-none on page 793](#)

- [privacy-password](#) on page 794
- [read-view](#) on page 795
- [remote-engine](#) on page 796
- [routing-instance](#) on page 797
- [security-level \(Defining Access Privileges\)](#) on page 798
- [security-level \(Generating SNMP Notifications\)](#) on page 799
- [security-model \(Access Privileges\)](#) on page 800
- [security-model \(Group\)](#) on page 801
- [security-model \(SNMP Notifications\)](#) on page 802
- [security-name \(Community String\)](#) on page 803
- [security-name \(Security Group\)](#) on page 804
- [security-name \(SNMP Notifications\)](#) on page 805
- [security-to-group](#) on page 806
- [snmp-community](#) on page 807
- [tag](#) on page 808
- [tag-list](#) on page 809
- [target-address](#) on page 810
- [target-parameters](#) on page 811
- [type](#) on page 812
- [user](#) on page 813
- [usm](#) on page 814
- [v3](#) on page 816
- [vacm](#) on page 819
- [write-view](#) on page 820

address

Supported Platforms	EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	<code>address <i>address</i>;</code>
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the SNMP target address.
Options	<i>address</i> —IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring the Address on page 138

address-mask

Supported Platforms	ACX Series , EX Series , M Series , MX Series , PTX Series , SRX Series , T Series
Syntax	<code>address-mask <i>address-mask</i>;</code>
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 on the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Verify the source addresses for a group of target addresses.
Options	<i>address-mask</i> combined with the address defines a range of addresses.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Configuring the Address Mask on page 138

authentication-key

Supported Platforms	M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	authentication-key <i>authentication-key</i> ;
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i> authentication-md5], [edit snmp v3 usm local-engine user <i>username</i> authentication-sha], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-md5], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> authentication-sha]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the key for user authentication, if you know it, instead of entering a cleartext password. The encrypted version of a password is known as a key.
Options	<i>authentication-key</i> —The encrypted version from the cleartext password.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MD5 Authentication on page 123• Configuring SHA Authentication on page 123

authentication-md5

Supported Platforms M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series

Syntax authentication-md5 {
 (authentication-key authentication-key | authentication-password
 authentication-password);
 }

Hierarchy Level [edit snmp v3 usm local-engine user *username*],
 [edit snmp v3 usm remote-engine *engine-id* user *username*]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure MD5 as the authentication type for the SNMPv3 user.



NOTE: You can only configure one authentication type for each SNMPv3 user.

For authentication, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
 snmp-control—To add this statement to the configuration.

Related Documentation • [Configuring MD5 Authentication on page 123](#)

authentication-none

Supported Platforms [ACX Series](#), [EX4600](#), [M Series](#), [MX Series](#), [OCX1100](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax authentication-none;

Hierarchy Level [edit snmp v3 usm local-engine user *username*],
[edit snmp v3 usm remote-engine *engine-id* user *username*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure that there should be no authentication for the SNMPv3 user.



NOTE: You can configure only one authentication type for each SNMPv3 user.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring No Authentication on page 124](#)

authentication-password

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax `authentication-password authentication-password;`

Hierarchy Level `[edit snmp v3 usm local-engine user username authentication-md5],`
`[edit snmp v3 usm local-engine user username authentication-sha],`
`[edit snmp v3 usm remote-engine engine-id user username authentication-md5],`
`[edit snmp v3 usm remote-engine engine-id user username authentication-sha]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the password for user authentication.

Options `authentication-password`—Password that a user enters. The password is then converted into a key that is used for authentication.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Required Privilege Level `snmp`—To view this statement in the configuration.
`snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring MD5 Authentication on page 123](#)
- [Configuring SHA Authentication on page 123](#)

authentication-sha

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax authentication-sha {
 ([authentication-key](#) *authentication-key* | [authentication-password](#)
 authentication-password);
}

Hierarchy Level [edit snmp v3 usm local-engine user *username*],
[edit snmp v3 usm remote-engine *engine-id* user *username*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the secure hash algorithm (SHA) as the authentication type for the SNMPv3 user.



NOTE: You can configure only one authentication type for each SNMPv3 user.

For authentication, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SHA Authentication on page 123](#)

community-name

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `community-name community-name;`

Hierarchy Level `[edit snmp v3 snmp-community community-index]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description The community name defines an SNMP community. The SNMP community authorizes SNMPv1 or SNMPv2 clients. The access privileges associated with the configured security name define which MIB objects are available and the operations (notify, read, or write) allowed on those objects.

Options *community-name*—Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose it in quotation marks (" ").



NOTE: Community names must be unique. You cannot configure the same community name at the `[edit snmp community]` and `[edit snmp v3 snmp-community community-index]` hierarchy levels.

The community name at the `[edit snmp v3 snmp-community community-index]` hierarchy level is encrypted and not displayed in the command-line interface (CLI).

Required Privilege Level `snmp`—To view this statement in the configuration.
`snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring the SNMPv3 Community on page 151](#)

context (SNMPv3)

Supported Platforms [M Series, MX Series, T Series](#)

Syntax `context context-name;`

Hierarchy Level [edit snmp v3 snmp-community *community-index*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the SNMPv3 context for access control. A context identifies a collection of information accessible for an SNMP entity.

Required Privilege snmp—To view this statement in the configuration.
Level snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the SNMPv3 Community on page 151](#)

engine-id

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `engine-id {
 (local engine-id-suffix | use-default-ip-address | use-mac-address);
}`

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.1 for EX Series switches.

Description The local engine ID is defined as the administratively unique identifier of an SNMPv3 engine, and is used for identification, not for addressing. There are two parts of an engine ID: prefix and suffix. The prefix is formatted according to the specifications defined in RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*. You can configure the suffix here.



NOTE: SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID. If you configure or change the engine ID, you must commit the new engine ID before you configure SNMPv3 users. Otherwise the keys generated from the configured passwords are based on the previous engine ID.

For the engine ID, we recommend using the MAC address of the management port.

Options `local engine-id-suffix`—Explicit setting for the engine ID suffix.

`use-default-ip-address`—The engine ID suffix is generated from the default IP address.

`use-mac-address`—The SNMP engine identifier is generated from the MAC address of the management interface on the router.

Default: `use-default-ip-address`

Required Privilege Level `snmp`—To view this statement in the configuration.
`snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring the Local Engine ID on page 120](#)

group (Configuring Group Name)

Supported Platforms ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series

Syntax

```
group group-name {  
  (default-context-prefix | context-prefix context-prefix){  
    security-model (any | usm | v1 | v2c) {  
      security-level (authentication | none | privacy) {  
        notify-view view-name;  
        read-view view-name;  
        write-view view-name;  
      }  
    }  
  }  
}
```

Hierarchy Level [edit snmp v3 vacm access]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The **default-context-prefix** statement, when included, adds all the contexts configured on the device to the group, whereas the **context-prefix context-prefix** statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series and OCX Series.) When the context prefix is specified as default (for example, **context-prefix default**), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as **logical system/routing instance**. For example, to specify routing instance ri1 in logical system ls1, include **context-prefix ls1/ri1**.

The remaining statements under this hierarchy are explained separately.

Options *group-name*—SNMPv3 group name created for the SNMPv3 group.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Group on page 127](#)

group (Defining Access Privileges for an SNMPv3 Group)

Supported Platforms	ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series
Syntax	<code>group <i>group-name</i>;</code>
Hierarchy Level	[edit snmp v3 vacm security-to-group security-model (usm v1 v2c) <code>security-name security-name</code>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Define access privileges granted to a group.
Options	<i>group-name</i> —Identifies a collection of SNMP security names that belong to the same access policy SNMP.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Group on page 132

retry-count

Supported Platforms [SRX Series](#)

Syntax `retry-count number;`

Hierarchy Level [edit snmp v3 [target-address](#) *target-address-name*]

Release Information Statement introduced in Junos OS Release 7.4.

Description Configure the retry count for SNMP informs.

Options *number*—Maximum number of times the inform is transmitted if no acknowledgment is received. If no acknowledgment is received after the inform is transmitted the maximum number of times, the inform message is discarded.

Default: 3 times

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Informs on page 143](#)
- [timeout on page 779](#)

timeout

Supported Platforms [SRX Series](#)

Syntax `timeout seconds;`

Hierarchy Level `[edit snmp v3 target-address target-address-name]`

Release Information Statement introduced in Junos OS Release 7.4.

Description Configure the timeout period (in seconds) for SNMP informs.

Options ***seconds***—Number of seconds to wait for an inform acknowledgment. If no acknowledgment is received within the timeout period, the inform is retransmitted.
Default: 15

Required Privilege Level `snmp`—To view this statement in the configuration.
 `snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP Informs on page 143](#)
- [retry-count on page 778](#)

local-engine

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax

```
local-engine {  
  user username {  
    authentication-md5 {  
      (authentication-key authentication-key | authentication-password  
        authentication-password);  
    }  
    authentication-none;  
    authentication-sha {  
      (authentication-key authentication-key | authentication-password  
        authentication-password);  
    }  
    privacy-aes128 {  
      (privacy-key privacy-key | privacy-password privacy-password);  
    }  
    privacy-des {  
      (privacy-key privacy-key | privacy-password privacy-password);  
    }  
    privacy-3des {  
      (privacy-key privacy-key | privacy-password privacy-password);  
    }  
    privacy-none {  
    }  
  }  
}
```

Hierarchy Level [edit snmp v3 [usm](#)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure local engine information for the user-based security model (USM).

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation • [Creating SNMPv3 Users on page 121](#)

message-processing-model

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series
Syntax	message-processing-model (v1 v2c v3);
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameter-name</i> parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the message processing model to be used when generating SNMP notifications.
Options	v1—SNMPv1 message process model. v2c—SNMPv2c message process model. v3—SNMPv3 message process model.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Message Processing Model on page 142

notify

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	<pre>notify <i>name</i> { tag <i>tag-name</i>; type (trap inform); }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>type inform option added in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.
Options	<p><i>name</i>—Name assigned to the notification.</p> <p><i>tag-name</i>—Notifications are sent to all targets configured with this tag.</p> <p><i>type</i>—Notification type is trap or inform. Traps are unconfirmed notifications. Informs are confirmed notifications.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Inform Notification Type and Target Address on page 148• Configuring the SNMPv3 Trap Notification on page 135

notify-filter (Applying to the Management Target)

Supported Platforms	ACX Series, EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	notify-filter <i>profile-name</i> ;
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameters-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify the notify filter to be used by a specific set of target parameters.
Options	<i>profile-name</i> —Name of the notify filter to apply to notifications.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Applying the Trap Notification Filter on page 141

notify-filter (Configuring the Profile Name)

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	<pre>notify-filter <i>profile-name</i> { oid <i>oid</i> (include exclude); }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.
Options	<p><i>profile-name</i>—Name assigned to the notify filter.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring the Trap Notification Filter on page 136• oid on page 786

notify-view

Supported Platforms	ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series
Syntax	notify-view <i>view-name</i> ;
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —Name of the view to which the SNMP user group has access.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring MIB Views on page 111 • Configuring the Notify View on page 129

oid

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	oid <i>oid</i> (include exclude);
Hierarchy Level	[edit snmp v3 notify-filter <i>profile-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.
Options	exclude —Exclude the subtree of MIB objects represented by the specified OID. include —Include the subtree of MIB objects represented by the specified OID. oid —Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Trap Notification Filter on page 136

parameters

Supported Platforms EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series

Syntax

```
parameters {
  message-processing-model (v1 | v2c | v3);
  security-level (none | authentication | privacy);
  security-model (usm | v1 | v2c);
  security-name security-name;
}
```

Hierarchy Level [edit snmp v3 target-parameters *target-parameters-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure a set of target parameters for message processing and security.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Defining and Configuring the Trap Target Parameters on page 140](#)

port

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax `port port-number;`

Hierarchy Level [edit snmp v3 target-address *target-address-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure a UDP port number for an SNMP target.

Default If you omit this statement, the default port is 162.

Options *port-number*—Port number for the SNMP target.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Port on page 138](#)

privacy-3des

Supported Platforms M Series, MX Series, QFX Series, SRX Series, T Series

Syntax `privacy-3des {
 (privacy-key privacy-key | privacy-password privacy-password);
}`

Hierarchy Level [edit snmp v3 usm local-engine user *username*],
[edit snmp v3 usm remote-engine *engine-id* user *username*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the triple Data Encryption Standard (3DES) as the privacy type for the SNMPv3 user.

For privacy encryption, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation • [Configuring the SNMPv3 Encryption Type on page 124](#)

privacy-aes128

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax `privacy-aes128 {
 (privacy-key privacy-key | privacy-password privacy-password);
}`

Hierarchy Level `[edit snmp v3 usm local-engine user username],
[edit snmp v3 usm remote-engine engine-id user username]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the Advanced Encryption Standard encryption algorithm (CFB128-AES-128 Privacy Protocol) for the SNMPv3 user.

For privacy encryption, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level `snmp`—To view this statement in the configuration.
`snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring the SNMPv3 Encryption Type on page 124](#)

privacy-des

Supported Platforms M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series

Syntax `privacy-des {
 (privacy-key privacy-key | privacy-password privacy-password);
}`

Hierarchy Level [edit snmp v3 usm local-engine user *username*],
[edit snmp v3 usm remote-engine *engine-id* user *username*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the Data Encryption Standard (DES) as the privacy type for the SNMPv3 user.

For privacy encryption, you can either enter a password in clear text, which is immediately encrypted, or, if you already have the encrypted version (known as a *key*), enter the key directly.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation • [Configuring the SNMPv3 Encryption Type on page 124](#)

privacy-key

Supported Platforms	M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	privacy-key <i>privacy-key</i> ;
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i> privacy-3des], [edit snmp v3 usm local-engine user <i>username</i> privacy-aes128], [edit snmp v3 usm local-engine user <i>username</i> privacy-des], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-3des], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-aes128], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i> privacy-des]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a privacy key for the SNMPv3 user, if you know it. The encrypted version of a password is known as a key.
Options	<i>privacy-key</i> —The encrypted version from the cleartext password. It is an alternative to entering a password.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMPv3 Encryption Type on page 124

privacy-none

Supported Platforms	M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	privacy-none;
Hierarchy Level	[edit snmp v3 usm local-engine user <i>username</i>], [edit snmp v3 usm remote-engine <i>engine-id</i> user <i>username</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure that no encryption be used for the SNMPv3 user.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMPv3 Encryption Type on page 124

privacy-password

Supported Platforms [M Series](#), [MX Series](#), [PTX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax `privacy-password privacy-password;`

Hierarchy Level `[edit snmp v3 usm local-engine user username privacy-3des],`
`[edit snmp v3 usm local-engine user username privacy-aes128],`
`[edit snmp v3 usm local-engine user username privacy-des],`
`[edit snmp v3 usm remote-engine engine-id user username privacy-3des],`
`[edit snmp v3 usm remote-engine engine-id user username privacy-aes128],`
`[edit snmp v3 usm remote-engine engine-id user username privacy-des]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure a privacy password for the SNMPv3 user.

Options *privacy-password*—Password that a user enters. The password is then converted into a key that is used for encryption.

SNMPv3 has special requirements when you create plain-text passwords on a router or switch:

- The password must be at least eight characters long.
- The password can include alphabetic, numeric, and special characters, but it cannot include control characters.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the SNMPv3 Encryption Type on page 124](#)

read-view

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	read-view <i>view-name</i> ;
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —The name of the view to which the SNMP user group has access.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Read View on page 130 • Configuring MIB Views on page 111

remote-engine

Supported Platforms EX4600, M Series, MX Series, OCX1100, PTX Series, QFX Series, SRX Series, T Series

Syntax

```
remote-engine engine-id {
  user username {
    authentication-md5 {
      (authentication-key authentication-key | authentication-password
        authentication-password);
    }
    authentication-none;
    authentication-sha {
      (authentication-key authentication-key | authentication-password
        authentication-password);
    }
    privacy-aes128 (privacy-key privacy-key | privacy-password privacy-password);
  }
  privacy-des {
    (privacy-key privacy-key | privacy-password privacy-password);
  }
  privacy-3des {
    (privacy-key privacy-key | privacy-password privacy-password);
  }
  privacy-none {
  }
}
```

Hierarchy Level [edit snmp v3 usm]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the remote engine information for the user-based security model (USM). To send inform messages to an SNMPv3 user on a remote device, you must configure the engine identifier for the SNMP agent on the remote device where the user resides.

Options *engine-id*—Specify engine identifier in hexadecimal format. Used to compute the security digest for authenticating and encrypting packets sent to a user on the remote host.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Remote Engine and Remote User on page 144](#)

routing-instance

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `routing-instance routing-instance-name;`

Hierarchy Level [edit snmp v3 [target-address](#) *target-address-name*]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify a routing instance for an SNMPv3 trap target.

Options *routing-instance-name*—Name of the routing instance.

To configure a routing instance within a logical system, specify the logical system name followed by the routing instance name. Use a slash (/) to separate the two names (for example, **test-ls/test-ri**). To configure the default routing instance on a logical system, specify the logical system name followed by **default** (for example, **test-ls/default**).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Trap Target Address on page 137](#)

security-level (Defining Access Privileges)

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series
Syntax	<pre>security-level (authentication none privacy) { notify-view view-name; read-view view-name; write-view view-name; }</pre>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the security level used for access privileges.
Default	none
Options	authentication —Provide authentication but no encryption. none —No authentication and no encryption. privacy —Provide authentication and encryption.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Security Level on page 128

security-level (Generating SNMP Notifications)

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series
Syntax	security-level (authentication none privacy);
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security level to use when generating SNMP notifications.
Default	none
Options	authentication —Provide authentication but no encryption. none —No authentication and no encryption. privacy —Provide authentication and encryption.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Security Level on page 142

security-model (Access Privileges)

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series
Syntax	security-model (usm v1 v2c);
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.
Options	usm —SNMPv3 security model. v1 —SNMPv1 security model. v2c —SNMPv2c security model.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Security Model on page 128

security-model (Group)

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#)

Syntax `security-model (usm | v1 | v2c) {
 security-name security-name {
 group group-name;
 }
}`

Hierarchy Level [edit snmp v3 vacm [security-to-group](#)]

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Define a security model for a group.

Options `usm`—SNMPv3 security model.
`v1`—SNMPv1 security model.
`v2c`—SNMPv2c security model.

Required Privilege `snmp`—To view this statement in the configuration.
Level `snmp-control`—To add this statement to the configuration.

Related Documentation

- [Configuring the Security Model on page 132](#)

security-model (SNMP Notifications)

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series
Syntax	security-model (usm v1 v2c);
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.
Options	usm —SNMPv3 security model. v1 —SNMPv1 security model. v2c —SNMPv2c security model.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Security Model on page 142

security-name (Community String)

Supported Platforms ACX Series, EX Series, M Series, MX Series, QFX Series, SRX Series, T Series

Syntax security-name *security-name*;

Hierarchy Level [edit snmp v3 *snmp-community* *community-index*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the [edit snmp v3 *snmp-community* *community-index*] hierarchy level.

Options *security-name*—Name that is used for messaging security and user access control.



NOTE: The security name must match the configured security name at the [edit snmp v3 target-parameters *target-parameters-name* parameters] hierarchy level when you configure traps or informs.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Security Names on page 152](#)

security-name (Security Group)

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [SRX Series](#), [T Series](#)

Syntax `security-name security-name {
 group group-name;
}`

Hierarchy Level `[edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c)]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Associate a group or a community string with a configured security group.

Options *security-name*—Username configured at the `[edit snmp v3 usm local-engine user username]` hierarchy level. For SNMPv1 and SNMPv2c, the security name is the community string configured at the `[edit snmp v3 snmp-community community-index]` hierarchy level.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Assigning Security Names to Groups on page 132](#)

security-name (SNMP Notifications)

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax `security-name security-name;`

Hierarchy Level [edit snmp v3 target-parameters *target-parameters-name* parameters]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the security name used when generating SNMP notifications.

Options *security-name*—If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification.



NOTE: The access privileges for the group associated with this security name must allow this notification to be sent.

If you are using the v1 or v2 security models, the security name at the [edit snmp v3 vacm security-to-group] hierarchy level must match the security name at the [edit snmp v3 snmp-community *community-index*] hierarchy level.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring the Security Name on page 143](#)

security-to-group

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [QFX Series](#), [SRX Series](#), [T Series](#)

Syntax

```
security-to-group {  
  security-model (usm | v1 | v2c) {  
    group group-name;  
    security-name security-name;  
  }  
}
```

Hierarchy Level [edit snmp v3 vacm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Assigning Security Model and Security Name to a Group on page 131](#)

snmp-community

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	<pre>snmp-community <i>community-index</i> { <i>community-name</i> <i>community-name</i>; <i>security-name</i> <i>security-name</i>; <i>tag</i> <i>tag-name</i>; }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the SNMP community.
Options	<i>community-index</i> —(Optional) String that identifies an SNMP community. The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the SNMPv3 Community on page 151

tag

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	tag <i>tag-name</i> ;
Hierarchy Level	[edit snmp v3 notify name], [edit snmp v3 snmp-community community-index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure a set of targets to receive traps or informs (for IPv4 packets only).
Options	<i>tag-name</i> —Identifies the address of managers that are allowed to use a community string.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Tag on page 152• Configuring the SNMPv3 Trap Notification on page 135

tag-list

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	tag-list <i>tag-list</i> ;
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an SNMP tag list used to select target addresses.
Options	<i>tag-list</i> —Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Trap Target Address on page 139

target-address

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax `target-address target-address-name {
 address address;
 address-mask address-mask;
 logical-system logical-system;
 port port-number;
 retry-count number;
 routing-instance instance;
 tag-list tag-list;
 target-parameters target-parameters-name;
 timeout seconds;
 }`

Hierarchy Level `[edit snmp v3]`

Release Information Statement introduced before Junos OS Release 7.4.
 Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the address of an SNMP management application and the parameters to be used in sending notifications.

Options ***target-address-name***—String that identifies the target address.

 The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level `snmp`—To view this statement in the configuration.
 `snmp-control`—To add this statement to the configuration.

Related Documentation • [Configuring the Trap Target Address on page 137](#)

target-parameters

Supported Platforms EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series

Syntax At the `[edit snmp v3]` hierarchy level:

```
target-parameters target-parameters-name {
  profile-name;
  parameters {
    message-processing-model (v1 | v2c | V3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
```

At the `[edit snmp v3 target-address target-address-name]` hierarchy level:

```
target-parameters target-parameters-name;
```

Hierarchy Level `[edit snmp v3]`
`[edit snmp v3 target-address target-address-name]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the `[edit snmp v3]` hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the `[edit snmp v3 target-parameters target-parameters-name]` hierarchy level to the target address configuration at the `[edit snmp v3]` hierarchy level.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Defining and Configuring the Trap Target Parameters on page 140](#)
- [Applying Target Parameters on page 139](#)

type

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	type (inform trap);
Hierarchy Level	[edit snmp v3 notify <i>name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>inform option added in Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Configure the type of SNMP notification.
Options	<p>inform—Defines the type of notification as an inform. SNMP informs are confirmed notifications.</p> <p>trap—Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP Informs on page 143• Configuring the SNMPv3 Trap Notification on page 135

user

Supported Platforms	M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series
Syntax	<code>user <i>username</i>;</code>
Hierarchy Level	[edit snmp v3 usm local-engine], [edit snmp v3 usm remote-engine <i>engine-id</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Specify a user associated with an SNMPv3 group on a local or remote SNMP engine.
Options	<i>username</i> —SNMPv3 user-based security model (USM) username.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating SNMPv3 Users on page 121

usm

Supported Platforms M Series, MX Series, PTX Series, QFX Series, SRX Series, T Series

```
Syntax  usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
        remote-engine engine-id {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-none;
                authentication-sha {
                    authentication-password authentication-password;
                }
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-3des {
                    privacy-password privacy-password;
                }
                privacy-none {
                    privacy-password privacy-password;
                }
            }
        }
    }
```


Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure user-based security model (USM) information. The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Creating SNMPv3 Users on page 121• Configuring the Remote Engine and Remote User on page 144

v3

Supported Platforms EX Series, M Series, MX Series, PTX Series, SRX Series, T Series

```

Syntax v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system logical-system;
        port port-number;
        retry-count number;
        routing-instance instance;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | V3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-none;
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-des {

```

```

        privacy-password privacy-password;
    }
    privacy-none;
}
}
remote-engine engine-id {
    user username {
        authentication-md5 {
            authentication-password authentication-password;
        }
        authentication-sha {
            authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
            privacy-password privacy-password;
        }
        privacy-des {
            privacy-password privacy-password;
        }
        privacy-3des {
            privacy-password privacy-password;
        }
        privacy-none {
            privacy-password privacy-password;
        }
    }
}
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
}
}

```

Hierarchy Level [edit snmp]

Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure SNMPv3. The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Minimum SNMPv3 Configuration on a Device Running Junos OS on page 116

vacm

Supported Platforms EX Series, M Series, MX Series, PTX Series, SRX Series, T Series

Syntax

```
vacm {
  access {
    group group-name {
      (default-context-prefix | context-prefix context-prefix){
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
  security-to-group {
    security-model (usm | v1 | v2c);
    security-name security-name {
      group group-name;
    }
  }
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure view-based access control model (VACM) information.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- [Defining Access Privileges for an SNMP Group on page 126](#)

write-view

Supported Platforms	ACX Series, EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, SRX Series, T Series
Syntax	<code>write-view view-name;</code>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series switches. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —Name of the view for which the SNMP user group has write permission.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MIB Views on page 111• Configuring the Write View on page 130

CHAPTER 44

Configuration Statements: TWAMP

- [twamp-server](#) on page 822
- [twamp-client](#) on page 824

twamp-server

Supported Platforms [SRX Series](#)

Syntax

```
services {  
  rpm {  
    twamp {  
      server {  
        authentication-mode;  
        max-connection-duration;  
        maximum-sessions;  
        maximum-sessions-per-connection;  
        maximum-connections;  
        maximum-connections-per-client;  
      }  
      port;  
      client-list abc {  
        address;  
      }  
    }  
  }  
}
```

Hierarchy Level [edit services rpm twamp server]

Release Information Statement introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Define the Two-Way Active Measurement Protocol (TWAMP) server configurations. This configuration describes a TWAMP server on the router that enables a TWAMP client to connect to the server using any media interface IP address. In such a scenario, the router functions as a TWAMP server and timestamping is performed in the ukernel of the media-facing Flexible PIC Concentrator (FPC).

Options **server**—TWAMP server configuration.

authentication-mode—The set of authentication modes the server can accept.

Default: none

max-connection-duration—The maximum time a connection can exist between the client and the server.

Default: 24

Range: 0-120

maximum-sessions—Maximum number of test sessions for the server.

Default: 64

Range: 1-2048

maximum-sessions-per-connection—Maximum number of test sessions per client connection.

Default: 64

Range: 1-1024

maximum-connections—Maximum number of connections for the server.

Default: 64

Range: 1-1000

maximum-connections-per-client—Maximum number of server connections per client.

Default: 64

Range: 1-500

port—TWAMP server listening port used by the routing instance.

Default: 862

Range: 1-65535

client-list—List of allowed control client hosts that can connect to this server. Each item in the list is a Classless Interdomain Routing (CIDR) address (ip address1/mask1) that represents the network of allowed hosts.

Required Privilege Level	system—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP</i>
------------------------------	--

twamp-client

Supported Platforms [SRX Series](#)

Syntax `control-connection control-connection-name {
 authentication-mode none;
 destination-port;
 history-size;
 moving-average-size;
 target-address;
 test-count ;
 test-interval;
 test-session test-session-name {
 target-address;
 data-fill-with-zeros;
 dscp-code-points;
 probe-count;
 probe-interval;
 thresholds;
 }
}`

Hierarchy Level [edit services rpm twamp client]

Release Information Statement introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Define the Two-Way Active Measurement Protocol (TWAMP) client configuration. A client opens a TCP connection to the server on well-known port 862. The host that initiates the TCP connection takes the roles of control client and session sender. The host that acknowledges the TCP connection accepts the roles of server and session reflector.

Options **client**—TWAMP client configuration.

authentication-mode—The set of authentication modes that the client can accept.

Default: none

destination-port—Either the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP) port to which a probe is sent.

Default: 862

history-size—Number of stored history entries.

Default: 50

Range: 0-500

moving-average-size—Number of samples for making statistical calculations.

Default: 0

Range: 0-1024

routing-instance—Routing instance used by test sessions.

target-address—Destination address or URL used for the probes.

test-count—Total number of test session iterations.

Default: 0

Range: 0-4,294,967,290

test-interval—Number of seconds to wait between tests.

Default: 1

Range: 1-255

test-session—Test session details.

data-fill-with-zeros—If the option is defined, then the test packet value becomes zero. If the option is not defined, then a pseudorandom number is updated as data.

dscp-code-points—DiffServ code point bits or alias used for UDP data of test probes or test packets in test sessions.

Default: 0

probe-count—Total number of probes per test.

Default: 1

Range: 1-4,294,967,290

thresholds—Test threshold values.

Required Privilege Level	interface—To view this statement in the configuration.
	interface-control—To add this statement to the configuration.

Related Documentation	<ul style="list-style-type: none">• <i>Two-Way Active Measurement Protocol Overview</i>
------------------------------	---

CHAPTER 45

Operational Commands

- clear chassis cluster ip-monitoring failure-count
- clear chassis cluster ip-monitoring failure-count ip-address
- clear ilmi statistics
- clear interfaces statistics
- clear services rpm twamp server connection
- clear snmp history
- clear snmp statistics
- request pppoe connect
- request pppoe disconnect
- request services ip-monitoring preempt-restore policy
- request services rpm twamp start
- request services rpm twamp stop
- request snmp spoof-trap
- show chassis alarms
- show chassis cluster ip-monitoring status redundancy-group
- show interfaces (SRX Series)
- show interfaces snmp-index
- show interfaces summary
- show ilmi statistics
- show security alarms
- show security datapath-debug capture
- show security datapath-debug counter
- show security monitoring
- show security monitoring fpc fpc-number
- show security monitoring performance session
- show security monitoring performance spu
- show services ip-monitoring status
- show services rpm twamp client connection

- `show services rpm twamp client history-results`
- `show services rpm twamp client probe-results`
- `show services rpm twamp client session`
- `show services rpm twamp server connection`
- `show services rpm twamp server session`
- `show snmp health-monitor`
- `show snmp inform-statistics`
- `show snmp mib`
- `show snmp rmon`
- `show snmp statistics`
- `show snmp stats-response-statistics`
- `show snmp v3`
- `show system alarms`
- `show system alarms`
- `show system resource-monitor fpc`

clear chassis cluster ip-monitoring failure-count

Supported Platforms [SRX Series, vSRX](#)

Syntax clear chassis cluster ip-monitoring failure-count

Release Information Command introduced in Junos OS Release 10.1.

Description Clear the failure count for all IP addresses.

Required Privilege Level clear

Related Documentation

- [clear chassis cluster ip-monitoring failure-count](#)
- [clear chassis cluster ip-monitoring failure-count ip-address on page 830](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> clear chassis cluster ip-monitoring failure-count
```

```
node0:
```

```
-----  
Cleared failure count for all IPs
```

```
node1:
```

```
-----  
Cleared failure count for all IPs
```

clear chassis cluster ip-monitoring failure-count ip-address

Supported Platforms [SRX Series, vSRX](#)

Syntax clear chassis cluster ip-monitoring failure-count ip-address 1.1.1.1

Release Information Command introduced in Junos OS Release 10.1.

Description Clear the failure count for a specified IP address.



NOTE: Entering an IP address at the end of this command is optional. If you do not specify an IP address, the failure count for all monitored IP addresses is cleared.

Required Privilege Level clear

Related Documentation

- [clear chassis cluster failover-count](#)
- [clear chassis cluster ip-monitoring failure-count on page 829](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
user@host> clear chassis cluster ip-monitoring failure-count ip-address 1.1.1.1

node0:
-----
Cleared failure count for IP: 1.1.1.1

node1:
-----
Cleared failure count for IP: 1.1.1.1
```


clear ilmi statistics

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Syntax clear ilmi statistics

Release Information Command introduced before Junos OS Release 7.4.

Description Set Integrated Local Management Interface (ILMI) statistics to zero.

Options This command has no options.

Required Privilege Level clear

Related Documentation

- [show ilmi statistics on page 888](#)

List of Sample Output [clear ilmi statistics on page 831](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ilmi statistics

```
user@host> clear ilmi statistics
```

clear interfaces statistics

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [SRX Series](#), [T Series](#)

Syntax clear interfaces statistics (all | *interface-name*)

Release Information Command introduced before Junos OS Release 7.4.

Description Set interface statistics to zero. If you issue the **clear interfaces statistics *interface-name*** command and then perform a graceful Routing Engine switchover, the interface statistics are not cleared on the new master. Reissue the command to clear the interface statistics again.

Starting in Junos OS Release 17.3R1, this command supports the clearing of Packet Forwarding Engine accounting statistics on logical interfaces configured with accounting options. On these interfaces, the current statistics values are stored as the new current baseline values and then the counters are reset to zero. If the **allow-clear** statement is included in the interface profile, then the cleared statistics values are reported to the accounting options flat file associated with the interface. Reporting is disabled by default; if **allow-clear** is not configured, then the CLI displays cleared statistics counters, but they are not reported to the flat file.

Options all—Set statistics on all interfaces to zero.

interface-name—Set statistics on a particular interface to zero.

Required Privilege Level clear

List of Sample Output [clear interfaces statistics on page 832](#)

Output Fields When you enter this command, you are provided no feedback on the status of your request.

Sample Output

clear interfaces statistics

```
user@host> clear interfaces statistics
```

clear services rpm twamp server connection

Supported Platforms [SRX Series](#)

Syntax `clear services rpm twamp server connection
<connection-id>`

Release Information Command introduced in Junos OS Release 15.1 for SRX Series devices.

Description Clear connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default all established connections are cleared (along with the sessions on those connections). To clear only a specific connection, specify the connection ID when you issue the command.

Options *connection-id*—(Optional) Clears specific connection as per the ID mentioned.

Required Privilege Level clear

clear snmp history

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFabric System, QFX Series, T Series
Syntax	clear snmp history (<i>index</i> all)
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Delete the history record of Simple Network Management Protocol (SNMP) samples of Ethernet statistics collected.
Options	all —Clear all the entries in the history index. index —Clear the contents of the specified entry in the history index.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• clear snmp statistics on page 835

clear snmp statistics

Supported Platforms	EX Series, M Series, MX Series, OCX1100, PTX Series, QFX Series, T Series
Syntax	clear snmp statistics
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
Description	Clear Simple Network Management Protocol (SNMP) statistics.
Options	This command has no options.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show snmp statistics on page 940
List of Sample Output	clear snmp statistics on page 835
Output Fields	See show snmp statistics for an explanation of output fields.

Sample Output

clear snmp statistics

In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too big: 0, No such names: 0, Bad values: 0,
    Read only: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0
  Output:
    Packets: 2298, Too big: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290

user@host> clear snmp statistics

```

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 0, Total set varbinds: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 0, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0
```

request pppoe connect

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX](#)

Syntax `request pppoe connect`

Release Information Statement supported on SRX300, SRX320, SRX340, and SRX345 is introduced in Junos OS Release 15.1X49-D60.
Statement supported on SRX1500 and vSRX instances is introduced in Junos OS Release 15.1X49-D100.

Description Connect all sessions that are down.

Options `pppoe interface name`— (Optional) Connect to a specified session.

Required Privilege Level maintenance

Related Documentation

- *Understanding PPPoE Interfaces*
- *Example: Configuring PPPoE Interfaces*

List of Sample Output [request pppoe connect on page 837](#)

Output Fields When you enter this command, this command returns no output.

Sample Output

`request pppoe connect`

```
user@host> request pppoe connect
```

request pppoe disconnect

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX](#)

Syntax `request pppoe disconnect`

Release Information Statement supported on SRX300, SRX320, SRX340, and SRX345 is introduced in Junos OS Release 15.1X49-D60.
Statement supported on SRX1500 and vSRX instances is introduced in Junos OS Release 15.1X49-D100.

Description Disconnect all active sessions.

Options **session id** — (Optional) Disconnect the session for which the session ID is specified.
pppoe interface name— (Optional) Disconnect the session for a specific pppoe interface name.

Required Privilege Level maintenance

Related Documentation

- [Understanding PPPoE Interfaces](#)
- [Example: Configuring PPPoE Interfaces](#)

List of Sample Output [request pppoe disconnect on page 838](#)

Output Fields When you enter this command, this command returns no output.

Sample Output

`request pppoe disconnect`

```
user@host> request pppoe disconnect
```


request services ip-monitoring preempt-restore policy

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX

Syntax request services ip-monitoring preempt-restore policy
<policy-name>

Release Information Command introduced in Junos OS Release 11.4.

Description If the no-preempt option is specified, the policy will not perform preemptive failback when it is in a failover state, and when the RPM probe test recovers from failure. To manually revert to the failback state, run the **request services ip-monitoring preempt-restore policy** command.



NOTE: The **request services ip-monitoring preempt-restore policy** command takes effect only when the RPM probe is in the pass state, and when the policy is in a failover state.

Options policy name—Name of the policy.

Required Privilege Level maintenance

Related Documentation • [show services ip-monitoring status on page 905](#)

List of Sample Output [run request services ip-monitoring preempt-restore policy <policy name> on page 839](#)

Output Fields When you run this command, the policy is restored to the failback state.

Sample Output

run request services ip-monitoring preempt-restore policy <policy name>

```
user@host> run request services ip-monitoring preempt-restore policy policy1
Restore request succeeded: Policy policy1
```

request services rpm twamp start

Supported Platforms [SRX Series](#)

Syntax `request services rpm twamp start client <control-connection-name>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Start the probes of a Two-Way Active Measurement Protocol (TWAMP) client. You can start all sessions for all TWAMP clients, or start a session for a specific TWAMP client. When you start test session configured for a particular TWAMP client, the control client initiates all requested testing with a start-sessions message, and the server sends an acknowledgment. If the control connection is not active between the server and the client, the control connection is also established and the test connections are started later. If the control client name is not specified, all the configured test sessions are commenced.

Options **start client**—Start the TWAMP session between the TWAMP client and the TWAMP server.

control-connection-name—(Optional) Start or stop the TWAMP session with the server only for the specified control connection or TWAMP control client.

Required Privilege Level view

List of Sample Output [request services rpm twamp start client on page 840](#)

Output Fields When you enter this command, you are not provided feedback on the status of your request.

Sample Output

`request services rpm twamp start client`

```
user@host> request services rpm twamp start client c1
```

request services rpm twamp stop

Supported Platforms [SRX Series](#)

Syntax `request services rpm twamp stop client <control-connection-name>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Stop the Two-Way Active Measurement Protocol (TWAMP) session between the TWAMP client and the TWAMP server. You can stop all sessions for all TWAMP clients, or stop a session for a specific TWAMP client. When you stop the test session, the control connection is closed only after the stop-sessions message is sent from the TWAMP client to the TWAMP server. If the control client name is not specified, all the configured test sessions are closed.

Options **stop client**—Stop the TWAMP session between the TWAMP client and the TWAMP server.
control-connection-name—(Optional) Start or stop the TWAMP session with the server only for the specified control connection or TWAMP control client.

Required Privilege Level view

List of Sample Output [request services rpm twamp stop client on page 841](#)

Output Fields When you enter this command, you are not provided feedback on the status of your request.

Sample Output

`request services rpm twamp stop client`

```
user@host> request services rpm twamp stop client c1
```

request snmp spoof-trap

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [OCX1100](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

Syntax `request snmp spoof-trap`
`<trap> variable-bindings <object> <instance> <value>`

Release Information Command introduced in Junos OS Release 8.2.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.

Options `<trap>`—Name of the trap to spoof.

`variable-bindings <object> <instance> <value>`—(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, `ifIndex[14] = 14`). Enclose the list of variable bindings in quotation marks (" ") and use a comma to separate each object name, instance, and value definition (for example, `variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"`). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.

`<dummy name>`—A dummy trap name to display the list of available traps.

`Question mark (?)`—Question mark? to display possible completions.

Required Privilege Level request

List of Sample Output [request snmp spoof-trap \(with Variable Bindings\) on page 842](#)
[request snmp spoof-trap \(Illegal Trap Name\) on page 842](#)
[request snmp spoof-trap \(Question Mark ?\) on page 847](#)

Sample Output

request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmp spoof-trap linkUp variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"
Spoof trap request result: trap sent successfully
```

request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found
```

```
Allowed Traps:
ads1AtucInitFailureTrap
```

adslAtucPerfESsThreshTrap
adslAtucPerfLofsThreshTrap
adslAtucPerfLolsThreshTrap
adslAtucPerfLossThreshTrap
adslAtucPerfLprsThreshTrap
adslAtucRateChangeTrap
adslAturPerfESsThreshTrap
adslAturPerfLofsThreshTrap
adslAturPerfLossThreshTrap
adslAturPerfLprsThreshTrap
adslAturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlswTrapCircuitDown
dlswTrapCircuitUp
dlswTrapTConnDown
dlswTrapTConnPartnerReject
dlswTrapTConnProtViolation
dlswTrapTConnUp
dsx1LineStatusChange
dsx3LineStatusChange
entConfigChange
fallingAlarm
frDLCIStatusChange
ggsnTrapChanged
ggsnTrapCleared
ggsnTrapNew
gmplsTunnelDown
ifMauJabberTrap
ipv6IfStateChange
isisAreaMismatch
isisAttemptToExceedMaxSequence
isisAuthenticationFailure
isisAuthenticationTypeFailure
isisCorruptedLSPDetected
isisDatabaseOverload
isisIDLenMismatch
isisLSPTooLargeToPropagate
isisManualAddressDrops
isisMaxAreaAddressesMismatch
isisOriginatingLSPBufferSizeMismatch
isisOwnLSPPurge
isisProtocolsSupportedMismatch
isisRejectedAdjacency
isisSequenceNumberSkip
isisVersionSkew
jnxAccessAuthServerDisabled
jnxAccessAuthServerEnabled
jnxAccessAuthServiceDown
jnxAccessAuthServiceUp
jnxBfdSessDetectionTimeHigh
jnxBfdSessTxIntervalHigh

jnxBgpM2BackwardTransition
jnxBgpM2Established
jnxCmCfgChange
jnxCmRescueChange
jnxCollFlowOverload
jnxCollFlowOverloadCleared
jnxCollFtpSwitchover
jnxCollMemoryAvailable
jnxCollMemoryUnavailable
jnxCollUnavailableDest
jnxCollUnavailableDestCleared
jnxCollUnsuccessfulTransfer
jnxDfcHardMemThresholdExceeded
jnxDfcHardMemUnderThreshold
jnxDfcHardPpsThresholdExceeded
jnxDfcHardPpsUnderThreshold
jnxDfcSoftMemThresholdExceeded
jnxDfcSoftMemUnderThreshold
jnxDfcSoftPpsThresholdExceeded
jnxDfcSoftPpsUnderThreshold
jnxEventTrap
jnxExampleStartup
jnxFEBSwitchover
jnxFanFailure
jnxFanOK
jnxFruCheck
jnxFruFailed
jnxFruInsertion
jnxFruOK
jnxFruOffline
jnxFruOnline
jnxFruPowerOff
jnxFruPowerOn
jnxFruRemoval
jnxHardDiskFailed
jnxHardDiskMissing
jnxJsAvPatternUpdateTrap
jnxJsChassisClusterSwitchover
jnxJsFwAuthCapacityExceeded
jnxJsFwAuthFailure
jnxJsFwAuthServiceDown
jnxJsFwAuthServiceUp
jnxJsNatAddrPoolThresholdStatus
jnxJsScreenAttack
jnxJsScreenCfgChange
jnxLdpLspDown
jnxLdpLspUp
jnxLdpSesDown
jnxLdpSesUp
jnxMIMstCistPortLoopProtectStateChangeTrap
jnxMIMstCistPortRootProtectStateChangeTrap
jnxMIMstErrTrap
jnxMIMstGenTrap
jnxMIMstInvalidBpduRxdTrap
jnxMIMstMstiPortLoopProtectStateChangeTrap
jnxMIMstMstiPortRootProtectStateChangeTrap
jnxMIMstNewRootTrap
jnxMIMstProtocolMigrationTrap
jnxMIMstRegionConfigChangeTrap
jnxMIMstTopologyChgTrap
jnxMacChangedNotification

jnxMplsLdpInitSesThresholdExceeded
jnxMplsLdpPathVectorLimitMismatch
jnxMplsLdpSessionDown
jnxMplsLdpSessionUp
jnxOspfV3IfConfigError
jnxOspfV3IfRxBadPacket
jnxOspfV3IfStateChange
jnxOspfV3LsdbApproachingOverflow
jnxOspfV3LsdbOverflow
jnxOspfV3NbrRestartHelperStatusChange
jnxOspfV3NbrStateChange
jnxOspfV3NssaTranslatorStatusChange
jnxOspfV3RestartStatusChange
jnxOspfV3VirtIfConfigError
jnxOspfV3VirtIfRxBadPacket
jnxOspfV3VirtIfStateChange
jnxOspfV3VirtNbrRestartHelperStatusChange
jnxOspfV3VirtNbrStateChange
jnxOtnAlarmCleared
jnxOtnAlarmSet
jnxOverTemperature
jnxPmonOverloadCleared
jnxPmonOverloadSet
jnxPingEgressJitterThresholdExceeded
jnxPingEgressStdDevThresholdExceeded
jnxPingEgressThresholdExceeded
jnxPingIngressJitterThresholdExceeded
jnxPingIngressStdDevThresholdExceeded
jnxPingIngressThresholdExceeded
jnxPingRttJitterThresholdExceeded
jnxPingRttStdDevThresholdExceeded
jnxPingRttThresholdExceeded
jnxPortBpduErrorStatusChangeTrap
jnxPortLoopProtectStateChangeTrap
jnxPortRootProtectStateChangeTrap
jnxPowerSupplyFailure
jnxPowerSupplyOK
jnxRedundancySwitchover
jnxRmonAlarmGetFailure
jnxRmonGetOk
jnxSecAccessIfMacLimitExceeded
jnxSecAccessDdsRateLimitCrossed
jnxSonetAlarmCleared
jnxSonetAlarmSet
jnxSpSvcSetCpuExceeded
jnxSpSvcSetCpuOk
jnxSpSvcSetZoneEntered
jnxSpSvcSetZoneExited
jnxStormEventNotification
jnxSyslogTrap
jnxTemperatureOK
jnxVccpPortDown
jnxVccpPortUp
jnxVpnIfDown
jnxVpnIfUp
jnxVpnPwDown
jnxVpnPwUp
jnx12aldGlobalMacLimit
jnx12aldInterfaceMacLimit
jnx12aldRoutingInstMacLimit
linkDown

linkUp
lldpRemTablesChange
mfrMibTrapBundleLinkMismatch
mplsLspChange
mplsLspDown
mplsLspInfoChange
mplsLspInfoDown
mplsLspInfoPathDown
mplsLspInfoPathUp
mplsLspInfoUp
mplsLspPathDown
mplsLspPathUp
mplsLspUp
mplsNumVrfRouteMaxThreshExceeded
mplsNumVrfRouteMidThreshExceeded
mplsNumVrfSecIllglLb1ThrshExcd
mplsTunnelDown
mplsTunnelReoptimized
mplsTunnelRerouted
mplsTunnelUp
mplsVrfIfDown
mplsVrfIfUp
mplsXCDown
mplsXCUp
msdpBackwardTransition
msdpEstablished
newRoot
ospfIfAuthFailure
ospfIfConfigError
ospfIfRxBadPacket
ospfIfStateChange
ospfLsdbApproachingOverflow
ospfLsdbOverflow
ospfMaxAgeLsa
ospfNbrStateChange
ospfOriginateLsa
ospfTxRetransmit
ospfVirtIfAuthFailure
ospfVirtIfConfigError
ospfVirtIfRxBadPacket
ospfVirtIfStateChange
ospfVirtIfTxRetransmit
ospfVirtNbrStateChange
pethMainPowerUsageOffNotification
pethMainPowerUsageOnNotification
pethPsePortOnOffNotification
pingProbeFailed
pingTestCompleted
pingTestFailed
ptopoConfigChange
risingAlarm
rpMauJabberTrap
sd1cLSStatusChange
sd1cPortStatusChange
topologyChange
traceRoutePathChange
traceRouteTestCompleted
traceRouteTestFailed
vrrpTrapAuthFailure
vrrpTrapNewMaster
warmStart

request snmp spoof-trap (Question Mark ?)

```

user@host> request snmp spoof-trap ?
Possible completions:
  <trap>                The name of the trap to spoof
  adslAtucInitFailureTrap
  adslAtucPerfESsThreshTrap
  adslAtucPerfLofsThreshTrap
  adslAtucPerfLolsThreshTrap
  adslAtucPerfLossThreshTrap
  adslAtucPerfLprsThreshTrap
  adslAtucRateChangeTrap
  adslAturPerfESsThreshTrap
  adslAturPerfLofsThreshTrap
  adslAturPerfLossThreshTrap
  adslAturPerfLprsThreshTrap
  adslAturRateChangeTrap
  apsEventChannelMismatch
  apsEventFEPLF
  apsEventModeMismatch
  apsEventPSBF
  apsEventSwitchover
  authenticationFailure
  bfdSessDown
  bfdSessUp
  bgpBackwardTransition
  bgpEstablished
  coldStart
  dlsWTrapCircuitDown
  dlsWTrapCircuitUp
---(more 10%)---

```

show chassis alarms

Supported Platforms [SRX Series](#)

Syntax `show chassis alarms`

Release Information Command introduced in Junos OS Release 11.1 for SRX Series devices.

Description Display information about the conditions that have been configured to trigger alarms.

Options This command has no options.

Additional Information Chassis alarms are preset. You cannot modify them.

You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm is lit, it indicates that you are running the device in a manner that we do not recommend.

On routers, you can manually silence external devices connected to the alarm relay contacts by pressing the alarm cutoff button, located on the craft interface. Silencing the device does not remove the alarm messages from the display (if present on the router) or extinguish the alarm LEDs. In addition, new alarms that occur after you silence an external device reactivate the external device.

In Junos OS Release 11.1 and later, alarms for fans also show the slot number of the fans in the CLI output.

Required Privilege Level view

Related Documentation

- [show system alarms on page 953](#)

List of Sample Output [show chassis alarms on page 849](#)

Output Fields [Table 135 on page 848](#) lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

Table 135: show chassis alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: Minor or Major.

Table 135: show chassis alarms Output Fields (*continued*)

Field Name	Field Description
Description	Information about the alarm.

Sample Output

show chassis alarms

```
user@host> show chassis alarms
4 alarms currently active
Alarm time           Class  Description
2012-05-29 16:47:18 UTC Major  /var partition usage crossed critical threshold
2012-05-29 16:47:18 UTC Minor  /var partition usage crossed high threshold
2012-05-29 16:47:18 UTC Major  /root partition usage crossed critical threshold
2012-05-29 16:47:18 UTC Minor  /root partition usage crossed high threshold
```

show chassis cluster ip-monitoring status redundancy-group

Supported Platforms [SRX Series, vSRX](#)

Syntax `show chassis cluster ip-monitoring status
<redundancy-group group-number>`

Release Information Command introduced in Junos OS Release 9.6. Support for global threshold, current threshold, and weight of each monitored IP address added in Junos OS Release 12.1X47-D10.

Description Display the status of all monitored IP addresses for a redundancy group.

- Options**
- `none`— Display the status of monitored IP addresses for all redundancy groups on the node.
 - `redundancy-group group-number`— Display the status of monitored IP addresses under the specified redundancy group.

Required Privilege Level view

Related Documentation

- [clear chassis cluster failover-count](#)

List of Sample Output [show chassis cluster ip-monitoring status on page 851](#)
[show chassis cluster ip-monitoring status redundancy-group on page 852](#)

Output Fields [Table 136 on page 850](#) lists the output fields for the `show chassis cluster ip-monitoring status` command.

Table 136: show chassis cluster ip-monitoring status Output Fields

Field Name	Field Description
Redundancy-group	ID number (0 - 255) of a redundancy group in the cluster.
Global threshold	Failover value for all IP addresses monitored by the redundancy group.
Current threshold	Value equal to the global threshold minus the total weight of the unreachable IP address.
IP Address	Monitored IP address in the redundancy group.
Status	Current reachability state of the monitored IP address. Values for this field are: reachable , unreachable , and unknown . The status is "unknown" if Packet Forwarding Engines (PFEs) are not yet up and running.
Failure count	Number of attempts to reach an IP address.

Table 136: show chassis cluster ip-monitoring status Output Fields (*continued*)

Field Name	Field Description
Reason	Explanation for the reported status. See Table 137 on page 851 .
Weight	Combined weight (0 - 255) assigned to all monitored IP addresses. A higher weight value indicates greater importance.

Expanded reason output fields for unreachable IP addresses added in Junos OS Release 10.1. You might see any of the following reasons displayed.

Table 137: show chassis cluster ip-monitoring status redundancy group Reason Fields

Reason	Reason Description
No route to host	The router could not resolve the ARP, which is needed to send the ICMP packet to the host with the monitored IP address.
No auxiliary IP found	The redundant Ethernet interface does not have an auxiliary IP address configured.
Reth child not up	A child interface of a redundant Ethernet interface is down.
redundancy-group state unknown	Unable to obtain the state (primary, secondary, secondary-hold, disable) of a redundancy-group.
No reth child MAC address	Could not extract the MAC address of the redundant Ethernet child interface.
Secondary link not monitored	The secondary link might be down (the secondary child interface of a redundant Ethernet interface is either down or non-functional).
Unknown	The IP address has just been configured and the router still does not know the status of this IP. or Do not know the exact reason for the failure.

Sample Output

show chassis cluster ip-monitoring status

```
user@host> show chassis cluster ip-monitoring status
node0:
```

```
-----
```

```
Redundancy group: 1
Global threshold: 200
Current threshold: -120
```

IP address	Status	Failure count	Reason	Weight
10.254.5.44	reachable	0	n/a	220
2.2.2.1	reachable	0	n/a	100

node1:

Redundancy group: 1
Global threshold: 200
Current threshold: -120

IP address	Status	Failure count	Reason	Weight
10.254.5.44	reachable	0	n/a	220
2.2.2.1	reachable	0	n/a	100

Sample Output

show chassis cluster ip-monitoring status redundancy-group

```
user@host> show chassis cluster ip-monitoring status redundancy-group 1
node0:
```

Redundancy group: 1

IP address	Status	Failure count	Reason
10.254.5.44	reachable	0	n/a
2.2.2.1	reachable	0	n/a
1.1.1.5	reachable	0	n/a
1.1.1.4	reachable	0	n/a
1.1.1.1	reachable	0	n/a

node1:

Redundancy group: 1

IP address	Status	Failure count	Reason
10.254.5.44	reachable	0	n/a
2.2.2.1	reachable	0	n/a
1.1.1.5	reachable	0	n/a
1.1.1.4	reachable	0	n/a
1.1.1.1	reachable	0	n/a

show interfaces (SRX Series)

Supported Platforms SRX Series, vSRX

Syntax show interfaces (
 <interface-name>
 <brief | detail | extensive | terse>
 <controller interface-name>|
 <descriptions interface-name>|
 <destination-class (all | destination-class-name logical-interface-name)>|
 <diagnostics optics interface-name>|
 <far-end-interval interface-fpc/pic/port>|
 <filters interface-name>|
 <flow-statistics interface-name>|
 <interval interface-name>|
 <load-balancing (detail | interface-name)>|
 <mac-database mac-address mac-address>|
 <mc-ae id identifier unit number revertive-info>|
 <media interface-name>|
 <policers interface-name>|
 <queue both-ingress-egress egress forwarding-class forwarding-class ingress l2-statistics>|
 <redundancy (detail | interface-name)>|
 <routing brief detail summary interface-name>|
 <routing-instance (all | instance-name)>|
 <snmp-index snmp-index>|
 <source-class (all | destination-class-name logical-interface-name)>|
 <statistics interface-name>|
 <switch-port switch-port number>|
 <transport pm (all | optics | otn) (all | current | currentday | interval | previousday) (all | interface-name)>|
 <zone interface-name>
)

Release Information Command modified in Junos OS Release 9.5.

Description Display status information and statistics about interfaces on SRX Series appliance running Junos OS.

On SRX Series appliance, on configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

- Options**
- **interface-name**—(Optional) Display standard information about the specified interface. Following is a list of typical interface names. Replace pim with the PIM slot and port with the port number.
 - **at-pim/0/port**—ATM-over-ADSL or ATM-over-SHDSL interface.
 - **ce1-pim/0/ port**—Channelized E1 interface.
 - **cl-0/0/8**—3G wireless modem interface for SRX320 devices.
 - **ct1-pim/0/port**—Channelized T1 interface.

- **dl0**—Dialer Interface for initiating ISDN and USB modem connections.
 - **e1-pim/0/port**—E1 interface.
 - **e3-pim/0/port**—E3 interface.
 - **fe-pim/0/port**—Fast Ethernet interface.
 - **ge-pim/0/port**—Gigabit Ethernet interface.
 - **se-pim/0/port**—Serial interface.
 - **t1-pim/0/port**—T1 (also called DS1) interface.
 - **t3-pim/0/port**—T3 (also called DS3) interface.
 - **wx-slot/0/0**—WAN acceleration interface, for the WXC Integrated Services Module (ISM 200).
-
- **brief | detail | extensive | terse**—(Optional) Display the specified level of output.
 - **controller**—(Optional) Show controller information.
 - **descriptions**—(Optional) Display interface description strings.
 - **destination-class**—(Optional) Show statistics for destination class.
 - **diagnostics**—(Optional) Show interface diagnostics information.
 - **far-end-interval**—(Optional) Show far end interval statistics.
 - **filters**—(Optional) Show interface filters information.
 - **flow-statistics**—(Optional) Show security flow counters and errors.
 - **interval**—(Optional) Show interval statistics.
 - **load-balancing**—(Optional) Show load-balancing status.
 - **mac-database**—(Optional) Show media access control database information.
 - **mc-ae**—(Optional) Show MC-AE configured interface information.
 - **media**—(Optional) Display media information.
 - **policers**—(Optional) Show interface policers information.
 - **queue**—(Optional) Show queue statistics for this interface.
 - **redundancy**—(Optional) Show redundancy status.
 - **routing**—(Optional) Show routing status.
 - **routing-instance**—(Optional) Name of routing instance.
 - **snmp-index**—(Optional) SNMP index of interface.
 - **source-class**—(Optional) Show statistics for source class.
 - **statistics**—(Optional) Display statistics and detailed output.
 - **switch-port**—(Optional) Front end port number (0..15).

- **transport**—(Optional) Show interface transport information.
- **zone**—(Optional) Interface's zone.

Required Privilege Level view

Related Documentation • *Understanding Layer 2 Interfaces*

List of Sample Output

- [show interfaces Gigabit Ethernet on page 863](#)
- [show interfaces brief \(Gigabit Ethernet\) on page 863](#)
- [show interfaces detail \(Gigabit Ethernet\) on page 864](#)
- [show interfaces statistics st0.0 detail on page 865](#)
- [show interfaces extensive \(Gigabit Ethernet\) on page 866](#)
- [show interfaces terse on page 869](#)
- [show interfaces controller \(Channelized E1 IQ with Logical E1\) on page 869](#)
- [show interfaces controller \(Channelized E1 IQ with Logical DSO\) on page 870](#)
- [show interfaces descriptions on page 870](#)
- [show interfaces destination-class all on page 870](#)
- [show interfaces diagnostics optics on page 870](#)
- [show interfaces far-end-interval coc12-5/2/0 on page 871](#)
- [show interfaces far-end-interval coc1-5/2/1:1 on page 872](#)
- [show interfaces filters on page 872](#)
- [show interfaces flow-statistics \(Gigabit Ethernet\) on page 872](#)
- [show interfaces interval \(Channelized OC12\) on page 873](#)
- [show interfaces interval \(E3\) on page 874](#)
- [show interfaces interval \(SONET/SDH\) on page 874](#)
- [show interfaces load-balancing on page 875](#)
- [show interfaces load-balancing detail on page 875](#)
- [show interfaces mac-database \(All MAC Addresses on a Port\) on page 875](#)
- [show interfaces mac-database \(All MAC Addresses on a Service\) on page 876](#)
- [show interfaces mac-database mac-address on page 876](#)
- [show interfaces mc-ae on page 877](#)
- [show interfaces media \(SONET/SDH\) on page 877](#)
- [show interfaces policers on page 877](#)
- [show interfaces policers interface-name on page 878](#)
- [show interfaces queue on page 878](#)
- [show interfaces redundancy on page 879](#)
- [show interfaces redundancy \(Aggregated Ethernet\) on page 879](#)
- [show interfaces redundancy detail on page 879](#)
- [show interfaces routing brief on page 879](#)
- [show interfaces routing detail on page 880](#)
- [show interfaces routing-instance all on page 880](#)
- [show interfaces snmp-index on page 881](#)
- [show interfaces source-class all on page 881](#)
- [show interfaces statistics \(Fast Ethernet\) on page 881](#)
- [show interfaces switch-port on page 882](#)
- [show interfaces transport pm on page 883](#)

[show security zones on page 884](#)

Output Fields [Table 138 on page 856](#) lists the output fields for the **show interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 138: show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Link-level type	Encapsulation being used on the physical interface.	All levels
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
MTU	Maximum transmission unit size on the physical interface.	All levels
Link mode	Link mode: Full-duplex or Half-duplex.	
Speed	Speed at which the interface is running.	All levels
BPDU error	Bridge protocol data unit (BPDU) error: Detected or None	
Loopback	Loopback status: Enabled or Disabled . If loopback is enabled, type of loopback: Local or Remote .	All levels
Source filtering	Source filtering status: Enabled or Disabled .	All levels
Flow control	Flow control status: Enabled or Disabled .	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: Enabled or Disabled .	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline. 	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the physical link.	All levels
CoS queues	Number of CoS queues configured.	detail extensive none

Table 138: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Current address	Configured MAC address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None
Output Rate	Output rate in bps and pps.	None
Active alarms and Active defects	<p>Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. These fields can contain the value None or Link.</p> <ul style="list-style-type: none"> • None—There are no active defects or alarms. • Link—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning. 	detail extensive none
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface. 	detail extensive

Table 138: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Input errors	<p>Input errors on the interface.</p> <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the ignore-l3-incompletes. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • FIFO errors—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • Resource errors—Sum of transmit drops. 	extensive
Output errors	<p>Output errors on the interface.</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Collisions—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation; therefore, for Gigabit Ethernet PICs, this number must always remain 0. If it is nonzero, there is a software bug. • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field must never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • FIFO errors—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the interfaces. • MTU errors—Number of packets whose size exceeded the MTU of the interface. • Resource errors—Sum of transmit drops. 	extensive

Table 138: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface.	extensive
Queue counters and queue number	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
MAC statistics	<p>Receive and Transmit statistics reported by the PIC's MAC subsystem, including the following:</p> <ul style="list-style-type: none"> • Total octets and total packets—Total number of octets and packets. • Unicast packets, Broadcast packets, and Multicast packets—Number of unicast, broadcast, and multicast packets. • CRC/Align errors—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • FIFO error—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • MAC control frames—Number of MAC control frames. • MAC pause frames—Number of MAC control frames with pause operational code. • Oversized frames—There are two possible conditions regarding the number of oversized frames: <ul style="list-style-type: none"> • Packet length exceeds 1518 octets, or • Packet length exceeds MRU • Jabber frames—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • Fragment frames—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets) and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • VLAN tagged frames—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • Code violations—Number of times an event caused the PHY to indicate "Data reception error" or "invalid data symbol error." 	extensive

Table 138: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Filter statistics	<p>Receive and Transmit statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected.</p> <ul style="list-style-type: none"> • Input packet count—Number of packets received from the MAC hardware that the filter processed. • Input packet rejects—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • Input DA rejects—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local device (which the router is rejecting). • Input SA rejects—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • Output packet count—Number of packets that the filter has given to the MAC hardware. • Output packet pad count—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • Output packet error count—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • CAM destination filters, CAM source filters—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields must be 0. 	extensive
Autonegotiation information	<p>Information about link autonegotiation.</p> <ul style="list-style-type: none"> • Negotiation status: <ul style="list-style-type: none"> • Incomplete—Ethernet interface has the speed or link mode configured. • No autonegotiation—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation. • Complete—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful. 	extensive
Packet Forwarding Engine configuration	<p>Information about the configuration of the Packet Forwarding Engine:</p> <ul style="list-style-type: none"> • Destination slot—FPC slot number. 	extensive

Table 138: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available. 	extensive
Interface transmit statistics	Status of the interface-transmit-statistics configuration: Enabled or Disabled.	detail extensive
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. 	detail extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP interface index number for the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • Input packets, Output packets—Number of packets received and transmitted on the interface set. 	detail extensive

Table 138: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Local statistics	Number and rate of bytes and packets destined to the device.	extensive
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the Output bytes and Output packets interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	extensive
Security	Security zones that interface belongs to.	extensive
Flow Input statistics	Statistics on packets received by flow module.	extensive
Flow Output statistics	Statistics on packets sent by flow module.	extensive
Flow error statistics (Packets dropped due to)	Statistics on errors in the flow module.	extensive
Protocol	Protocol family.	detail extensive none
MTU	Maximum transmission unit size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route Table	Route table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	detail extensive none
Flags	Information about protocol family flags. .	detail extensive
Addresses, Flags	Information about the address flags..	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address of the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces Gigabit Ethernet

```

user@host> show interfaces ge-0/0/1
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:42 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Input packets : 0
  Output packets: 0
  Security: Zone: public
  Protocol inet, MTU: 1500
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255

```

Sample Output

show interfaces brief (Gigabit Ethernet)

```

user@host> show interfaces ge-3/0/2 brief
Physical interface: ge-3/0/2, Enabled, Physical link is Up
  Link-level type: 52, MTU: 1522, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface ge-3/0/2.0
  Flags: SNMP-Traps 0x4000
  VLAN-Tag [ 0x8100.512 0x8100.513 ] In(pop-swap 0x8100.530) Out(swap-push
0x8100.512 0x8100.513)
  Encapsulation: VLAN-CCC
  ccc

Logical interface ge-3/0/2.32767
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x0000.0 ] Encapsulation: ENET2

```

Sample Output

show interfaces detail (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1 detail
Physical interface: ge-0/0/1, Enabled, Physical link is Down
  Interface index: 135, SNMP ifIndex: 510, Generation: 138
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled, Source filtering:
  Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues    : 8 supported, 8 maximum usable queues
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
  Last flapped   : 2015-05-12 08:36:59 UTC (1w2d 00:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
  Egress queues: 8 supported, 4 in use
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort           0                0                0
    1 expedited-fo         0                0                0
    2 assured-forw         0                0                0
    3 network-cont         0                0                0

  Queue number:      Mapped forwarding classes
    0                best-effort
    1                expedited-forwarding
    2                assured-forwarding
    3                network-control
  Active alarms  : LINK
  Active defects : LINK
  Interface transmit statistics: Disabled

  Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
    Input bytes  :                0
    Output bytes :                0
    Input packets:                0
    Output packets:              0
  Local statistics:
    Input bytes  :                0
    Output bytes :                0
    Input packets:                0
    Output packets:              0
  Transit statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps

```

```

Output packets:                                0                0 pps
Security: Zone: public
Flow Statistics :
Flow Input statistics :
  Self packets :                               0
  ICMP packets :                               0
  VPN packets :                                0
  Multicast packets :                          0
  Bytes permitted by policy :                   0
  Connections established :                     0
Flow Output statistics:
  Multicast packets :                           0
  Bytes permitted by policy :                     0
Flow error statistics (Packets dropped due to):
  Address spoofing:                             0
  Authentication failed:                         0
  Incoming NAT errors:                          0
  Invalid zone received packet:                  0
  Multiple user authentications:                 0
  Multiple incoming NAT:                        0
  No parent for a gate:                         0
  No one interested in self packets:             0
  No minor session:                             0
  No more sessions:                             0
  No NAT gate:                                  0
  No route present:                             0
  No SA for incoming SPI:                       0
  No tunnel found:                              0
  No session for a gate:                        0
  No zone or NULL zone binding                  0
  Policy denied:                                0
  Security association not active:               0
  TCP sequence number out of window:            0
  Syn-attack protection:                        0
  User authentication errors:                   0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
Flags: Sendbcast-pkt-to-re
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255, Generation:
150

```

show interfaces statistics st0.0 detail

```

user@host> show interfaces statistics st0.0 detail
Logical interface st0.0 (Index 71) (SNMP ifIndex 609) (Generation 136)
Flags: Up Point-To-Point SNMP-Traps Encapsulation: Secure-Tunnel
Traffic statistics:
  Input bytes :                               528152756774
  Output bytes :                              575950643520
  Input packets:                             11481581669
  Output packets:                             12520666095
Local statistics:
  Input bytes :                               0
  Output bytes :                              0
  Input packets:                             0
  Output packets:                             0
Transit statistics:
  Input bytes :                               0                121859888 bps
  Output bytes :                              0                128104112 bps
  Input packets:                             0                331141 pps
  Output packets:                             0                348108 pps

```

```

Security: Zone: untrust
Allowed host-inbound traffic : any-service bfd bgp dvmrp igmp ldp msdp nhrp
ospf ospf3 pgm pim rip ripng router-discovery rsvp
sap vrrp
Flow Statistics :
Flow Input statistics :
  Self packets : 0
  ICMP packets : 0
  VPN packets : 0
  Multicast packets : 0
  Bytes permitted by policy : 525984295844
  Connections established : 7
Flow Output statistics:
  Multicast packets : 0
  Bytes permitted by policy : 576003290222
Flow error statistics (Packets dropped due to):
  Address spoofing: 0
  Authentication failed: 0
  Incoming NAT errors: 0
  Invalid zone received packet: 0
  Multiple user authentications: 0
  Multiple incoming NAT: 0
  No parent for a gate: 0
  No one interested in self packets: 0
  No minor session: 0
  No more sessions: 0
  No NAT gate: 0
  No route present: 2000280
  No SA for incoming SPI: 0
  No tunnel found: 0
  No session for a gate: 0
  No zone or NULL zone binding 0
  Policy denied: 0
  Security association not active: 0
  TCP sequence number out of window: 0
  Syn-attack protection: 0
  User authentication errors: 0
Protocol inet, MTU: 9192
Max nh cache: 0, New hold nh limit: 0, Curr nh cnt: 0, Curr new hold cnt: 0,
NH drop cnt: 0
Generation: 155, Route table: 0
Flags: Sendbroadcast-pkt-to-re

```

Sample Output

show interfaces extensive (Gigabit Ethernet)

```

user@host> show interfaces ge-0/0/1.0 extensive
Physical interface: ge-0/0/1, Enabled, Physical link is Down
Interface index: 135, SNMP ifIndex: 510, Generation: 138
Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed: 1000mbps,

BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
Device flags : Present Running Down
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Link flags : None
CoS queues : 8 supported, 8 maximum usable queues
Hold-times : Up 0 ms, Down 0 ms

```

```

Current address: 00:1f:12:e4:b1:01, Hardware address: 00:1f:12:e4:b1:01
Last flapped   : 2015-05-12 08:36:59 UTC (1w1d 22:57 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes   :                0                0 bps
Output bytes  :                0                0 bps
Input packets :                0                0 pps
Output packets:                0                0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Policed discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
FIFO errors: 0, Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 best-effort	0	0	0
1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

```

Queue number:      Mapped forwarding classes
0                  best-effort
1                  expedited-forwarding
2                  assured-forwarding
3                  network-control
Active alarms : LINK
Active defects : LINK
MAC statistics:

```

	Receive	Transmit
Total octets	0	0
Total packets	0	0
Unicast packets	0	0
Broadcast packets	0	0
Multicast packets	0	0
CRC/Align errors	0	0
FIFO errors	0	0
MAC control frames	0	0
MAC pause frames	0	0
Oversized frames	0	0
Jabber frames	0	0
Fragment frames	0	0
VLAN tagged frames	0	0
Code violations	0	0

```

Filter statistics:
Input packet count      0
Input packet rejects    0
Input DA rejects        0
Input SA rejects        0
Output packet count     0
Output packet pad count 0
Output packet error count 0
CAM destination filters: 2, CAM source filters: 0
Autonegotiation information:
Negotiation status: Incomplete
Packet Forwarding Engine configuration:

```

```

    Destination slot: 0
    CoS information:
    Direction : Output
    CoS transmit queue
Limit      Bandwidth      Buffer Priority
            %      bps      %      usec
    0 best-effort      95      950000000      95      0      low
none
    3 network-control      5      50000000      5      0      low
none

```

Interface transmit statistics: Disabled

Logical interface ge-0/0/1.0 (Index 71) (SNMP ifIndex 514) (Generation 136)

Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2

Traffic statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Local statistics:

```

Input bytes :      0
Output bytes :      0
Input packets:      0
Output packets:      0

```

Transit statistics:

```

Input bytes :      0      0 bps
Output bytes :      0      0 bps
Input packets:      0      0 pps
Output packets:      0      0 pps

```

Security: Zone: public

Flow Statistics :

Flow Input statistics :

```

Self packets :      0
ICMP packets :      0
VPN packets :      0
Multicast packets :      0
Bytes permitted by policy :      0
Connections established :      0

```

Flow Output statistics:

```

Multicast packets :      0
Bytes permitted by policy :      0

```

Flow error statistics (Packets dropped due to):

```

Address spoofing:      0
Authentication failed:      0
Incoming NAT errors:      0
Invalid zone received packet:      0
Multiple user authentications:      0
Multiple incoming NAT:      0
No parent for a gate:      0
No one interested in self packets:      0
No minor session:      0
No more sessions:      0
No NAT gate:      0
No route present:      0
No SA for incoming SPI:      0
No tunnel found:      0
No session for a gate:      0
No zone or NULL zone binding      0
Policy denied:      0
Security association not active:      0
TCP sequence number out of window:      0

```

```

Syn-attack protection:          0
User authentication errors:     0
Protocol inet, MTU: 1500, Generation: 150, Route table: 0
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.1, Broadcast: 1.1.1.255,
Generation: 150

```

Sample Output

show interfaces terse

```

user@host> show interfaces terse

```

Interface	Admin	Link	Proto	Local	Remote
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet	10.209.4.61/18	
gr-0/0/0	up	up			
ip-0/0/0	up	up			
st0	up	up			
st0.1	up	ready	inet		
ls-0/0/0	up	up			
lt-0/0/0	up	up			
mt-0/0/0	up	up			
pd-0/0/0	up	up			
pe-0/0/0	up	up			
e3-1/0/0	up	up			
t3-2/0/0	up	up			
e1-3/0/0	up	up			
se-4/0/0	up	down			
t1-5/0/0	up	up			
br-6/0/0	up	up			
dc-6/0/0	up	up			
dc-6/0/0.32767	up	up			
bc-6/0/0:1	down	up			
bc-6/0/0:1.0	up	down			
d10	up	up			
d10.0	up	up	inet		
dsc	up	up			
gre	up	up			
ipip	up	up			
lo0	up	up			
lo0.16385	up	up	inet	10.0.0.1 10.0.0.16	--> 0/0 --> 0/0
lsi	up	up			
mtun	up	up			
pimd	up	up			
pime	up	up			
pp0	up	up			

Sample Output

show interfaces controller (Channelized E1 IQ with Logical E1)

```

user@host> show interfaces controller ce1-1/2/6

```

Controller	Admin	Link
ce1-1/2/6	up	up
e1-1/2/6	up	up

show interfaces controller (Channelized E1 IQ with Logical DSO)

```
user@host> show interfaces controller ce1-1/2/3
```

Controller	Admin	Link
ce1-1/2/3	up	up
ds-1/2/3:1	up	up
ds-1/2/3:2	up	up

Sample Output**show interfaces descriptions**

```
user@host> show interfaces descriptions
```

Interface	Admin	Link	Description
so-1/0/0	up	up	M20-3#1
so-2/0/0	up	up	GSR-12#1
ge-3/0/0	up	up	SMB-OSPF_Area300
so-3/3/0	up	up	GSR-13#1
so-3/3/1	up	up	GSR-13#2
ge-4/0/0	up	up	T320-7#1
ge-5/0/0	up	up	T320-7#2
so-7/1/0	up	up	M160-6#1
ge-8/0/0	up	up	T320-7#3
ge-9/0/0	up	up	T320-7#4
so-10/0/0	up	up	M160-6#2
so-13/0/0	up	up	M20-3#2
so-14/0/0	up	up	GSR-12#2
ge-15/0/0	up	up	SMB-OSPF_Area100
ge-15/0/1	up	up	GSR-13#3

Sample Output**show interfaces destination-class all**

```
user@host> show interfaces destination-class all
```

```
Logical interface so-4/0/0.0
```

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	0	0
	(0)	(0)
silver	0	0
	(0)	(0)

```
Logical interface so-0/1/3.0
```

Destination class	Packets (packet-per-second)	Bytes (bits-per-second)
gold	0	0
	(0)	(0)
silver	0	0
	(0)	(0)

Sample Output**show interfaces diagnostics optics**

```
user@host> show interfaces diagnostics optics ge-2/0/0
```



```

Physical interface: ge-2/0/0
Laser bias current           : 7.408 mA
Laser output power          : 0.3500 mW / -4.56 dBm
Module temperature          : 23 degrees C / 73 degrees F
Module voltage              : 3.3450 V
Receiver signal average optical power : 0.0002 mW / -36.99 dBm
Laser bias current high alarm : Off
Laser bias current low alarm  : Off
Laser bias current high warning : Off
Laser bias current low warning : Off
Laser output power high alarm : Off
Laser output power low alarm  : Off
Laser output power high warning : Off
Laser output power low warning : Off
Module temperature high alarm : Off
Module temperature low alarm  : Off
Module temperature high warning : Off
Module temperature low warning : Off
Module voltage high alarm     : Off
Module voltage low alarm      : Off
Module voltage high warning   : Off
Module voltage low warning    : Off
Laser rx power high alarm     : Off
Laser rx power low alarm      : On
Laser rx power high warning   : Off
Laser rx power low warning    : On
Laser bias current high alarm threshold : 17.000 mA
Laser bias current low alarm threshold : 1.000 mA
Laser bias current high warning threshold : 14.000 mA
Laser bias current low warning threshold : 2.000 mA
Laser output power high alarm threshold : 0.6310 mW / -2.00 dBm
Laser output power low alarm threshold : 0.0670 mW / -11.74 dBm
Laser output power high warning threshold : 0.6310 mW / -2.00 dBm
Laser output power low warning threshold : 0.0790 mW / -11.02 dBm
Module temperature high alarm threshold : 95 degrees C / 203 degrees F
Module temperature low alarm threshold : -25 degrees C / -13 degrees F
Module temperature high warning threshold : 90 degrees C / 194 degrees F
Module temperature low warning threshold : -20 degrees C / -4 degrees F
Module voltage high alarm threshold : 3.900 V
Module voltage low alarm threshold : 2.700 V
Module voltage high warning threshold : 3.700 V
Module voltage low warning threshold : 2.900 V
Laser rx power high alarm threshold : 1.2590 mW / 1.00 dBm
Laser rx power low alarm threshold : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold : 0.7940 mW / -1.00 dBm
Laser rx power low warning threshold : 0.0158 mW / -18.01 dBm

```

Sample Output

show interfaces far-end-interval coc12-5/2/0

```

user@host> show interfaces far-end-interval coc12-5/2/0
Physical interface: coc12-5/2/0, SNMP ifIndex: 121
05:30-current:
  ES-L: 1, SES-L: 1, UAS-L: 0
05:15-05:30:
  ES-L: 0, SES-L: 0, UAS-L: 0
05:00-05:15:
  ES-L: 0, SES-L: 0, UAS-L: 0
04:45-05:00:

```

```

    ES-L: 0, SES-L: 0, UAS-L: 0
04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0
04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0
04:00-04:15:
...

```

show interfaces far-end-interval coc1-5/2/1:1

```

user@host> run show interfaces far-end-interval coc1-5/2/1:1
Physical interface: coc1-5/2/1:1, SNMP ifIndex: 342
05:30-current:
    ES-L: 1, SES-L: 1, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:15-05:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
05:00-05:15:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:45-05:00:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:30-04:45:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:15-04:30:
    ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0, SES-P: 0, UAS-P: 0
04:00-04:15:

```

Sample Output

show interfaces filters

```

user@host> show interfaces filters

```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/0/0	up	up			
ge-0/0/0.0	up	up	inet		
			iso		
ge-5/0/0	up	up			
ge-5/0/0.0	up	up	any		f-any
			inet		f-inet
			multiservice		
gr-0/3/0	up	up			
ip-0/3/0	up	up			
mt-0/3/0	up	up			
pd-0/3/0	up	up			
pe-0/3/0	up	up			
vt-0/3/0	up	up			
at-1/0/0	up	up			
at-1/0/0.0	up	up	inet		
			iso		
at-1/1/0	up	down			
at-1/1/0.0	up	down	inet		
			iso		
....					

Sample Output

show interfaces flow-statistics (Gigabit Ethernet)

```

user@host> show interfaces flow-statistics ge-0/0/1.0

```

```

Logical interface ge-0/0/1.0 (Index 70) (SNMP ifIndex 49)
  Flags: SNMP-Traps Encapsulation: ENET2
  Input packets : 5161
  Output packets: 83
  Security: Zone: zone2
  Allowed host-inbound traffic : bootp bfd bgp dns dvmrp ldp msdp nhrp ospf
pgm
  pim rip router-discovery rsvp sap vrrp dhcp finger ftp tftp ident-reset http
https ike
  netconf ping rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
xnm-ssl
  lsping
  Flow Statistics :
  Flow Input statistics :
    Self packets : 0
    ICMP packets : 0
    VPN packets : 2564
    Bytes permitted by policy : 3478
    Connections established : 1
  Flow Output statistics:
    Multicast packets : 0
    Bytes permitted by policy : 16994
  Flow error statistics (Packets dropped due to):
    Address spoofing: 0
    Authentication failed: 0
    Incoming NAT errors: 0
    Invalid zone received packet: 0
    Multiple user authentications: 0
    Multiple incoming NAT: 0
    No parent for a gate: 0
    No one interested in self packets: 0
    No minor session: 0
    No more sessions: 0
    No NAT gate: 0
    No route present: 0
    No SA for incoming SPI: 0
    No tunnel found: 0
    No session for a gate: 0
    No zone or NULL zone binding 0
    Policy denied: 0
    Security association not active: 0
    TCP sequence number out of window: 0
    Syn-attack protection: 0
    User authentication errors: 0
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113.1/24, Local: 203.0.113.2, Broadcast: 2.2.2.255

```

Sample Output

show interfaces interval (Channelized OC12)

```

user@host> show interfaces interval t3-0/3/0:0
Physical interface: t3-0/3/0:0, SNMP ifIndex: 23
17:43-current:
  LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
  SEFS: 0, UAS: 0
17:28-17:43:
  LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,

```

```

    SEFS: 0, UAS: 0
17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
...
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

show interfaces interval (E3)

```

user@host> show interfaces interval e3-0/3/0
Physical interface: e3-0/3/0, SNMP ifIndex: 23
17:43-current:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
17:28-17:43:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
17:13-17:28:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
16:58-17:13:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
    SEFS: 0, UAS: 0
16:43-16:58:
    LCV: 0, PCV: 0, CCV: 0, LES: 0, PES: 0, PSES: 0, CES: 0, CSES: 0,
....
Interval Total:
    LCV: 230, PCV: 1145859, CCV: 455470, LES: 0, PES: 230, PSES: 230,
    CES: 230, CSES: 230, SEFS: 230, UAS: 238

```

show interfaces interval (SONET/SDH)

```

user@host> show interfaces interval so-0/1/0
Physical interface: so-0/1/0, SNMP ifIndex: 19
20:02-current:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
19:47-20:02:
    ES-S: 267, SES-S: 267, SEFS-S: 267, ES-L: 267, SES-L: 267, UAS-L: 267,
    ES-P: 267, SES-P: 267, UAS-P: 267
19:32-19:47:
    ES-S: 56, SES-S: 56, SEFS-S: 56, ES-L: 56, SES-L: 56, UAS-L: 46, ES-P: 56,
    SES-P: 56, UAS-P: 46
19:17-19:32:
    ES-S: 0, SES-S: 0, SEFS-S: 0, ES-L: 0, SES-L: 0, UAS-L: 0, ES-P: 0,
    SES-P: 0, UAS-P: 0
19:02-19:17:
.....

```

Sample Output

show interfaces load-balancing

```
user@host> show interfaces load-balancing
Interface  State           Last change  Member count
ams0       Up              1d 00:50    2
ams1       Up              00:00:59    2
```

show interfaces load-balancing detail

```
user@host> show interfaces load-balancing detail
Load-balancing interfaces detail
Interface   : ams0
State       : Up
Last change : 1d 00:51
Member count : 2
Members     :
  Interface  Weight  State
  mams-2/0/0  10     Active
  mams-2/1/0  10     Active
```

Sample Output

show interfaces mac-database (All MAC Addresses on a Port)

```
user@host> show interfaces mac-database xe-0/3/3
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
MAC address      Input frames  Input bytes  Output frames  Output bytes
00:00:00:00:00:00      1           56           0             0
00:00:c0:01:01:02     7023810     323095260    0             0
00:00:c0:01:01:03     7023810     323095260    0             0
00:00:c0:01:01:04     7023810     323095260    0             0
00:00:c0:01:01:05     7023810     323095260    0             0
00:00:c0:01:01:06     7023810     323095260    0             0
00:00:c0:01:01:07     7023810     323095260    0             0
00:00:c0:01:01:08     7023809     323095214    0             0
00:00:c0:01:01:09     7023809     323095214    0             0
00:00:c0:01:01:0a     7023809     323095214    0             0
00:00:c0:01:01:0b     7023809     323095214    0             0
00:00:c8:01:01:02     30424784    1399540064    37448598      1722635508
00:00:c8:01:01:03     30424784    1399540064    37448598      1722635508
00:00:c8:01:01:04     30424716    1399536936    37448523      1722632058
00:00:c8:01:01:05     30424789    1399540294    37448598      1722635508
00:00:c8:01:01:06     30424788    1399540248    37448597      1722635462
00:00:c8:01:01:07     30424783    1399540018    37448597      1722635462
00:00:c8:01:01:08     30424783    1399540018    37448596      1722635416
00:00:c8:01:01:09     8836796     406492616    8836795      406492570
```

```

00:00:c8:01:01:0a      30424712      1399536752      37448521      1722631966
00:00:c8:01:01:0b      30424715      1399536890      37448523      1722632058
Number of MAC addresses : 21

```

show interfaces mac-database (All MAC Addresses on a Service)

```

user@host> show interfaces mac-database xe-0/3/3
Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address      Input frames   Input bytes   Output frames   Output bytes
00:00:00:00:00:00      1             56            0              0
00:00:c0:01:01:02     7023810      323095260     0              0
00:00:c0:01:01:03     7023810      323095260     0              0
00:00:c0:01:01:04     7023810      323095260     0              0
00:00:c0:01:01:05     7023810      323095260     0              0
00:00:c0:01:01:06     7023810      323095260     0              0
00:00:c0:01:01:07     7023810      323095260     0              0
00:00:c0:01:01:08     7023809      323095214     0              0
00:00:c0:01:01:09     7023809      323095214     0              0
00:00:c0:01:01:0a     7023809      323095214     0              0
00:00:c0:01:01:0b     7023809      323095214     0              0
00:00:c8:01:01:02     31016568     1426762128    38040381      1749857526
00:00:c8:01:01:03     31016568     1426762128    38040382      1749857572
00:00:c8:01:01:04     31016499     1426758954    38040306      1749854076
00:00:c8:01:01:05     31016573     1426762358    38040381      1749857526
00:00:c8:01:01:06     31016573     1426762358    38040381      1749857526
00:00:c8:01:01:07     31016567     1426762082    38040380      1749857480
00:00:c8:01:01:08     31016567     1426762082    38040379      1749857434
00:00:c8:01:01:09     9428580      433714680     9428580       433714680
00:00:c8:01:01:0a     31016496     1426758816    38040304      1749853984
00:00:c8:01:01:0b     31016498     1426758908    38040307      1749854122

```

show interfaces mac-database mac-address

```

user@host> show interfaces mac-database xe-0/3/3 mac-address 00:00:c8:01:01:09
Physical interface: xe-0/3/3, Enabled, Physical link is Up
  Interface index: 372, SNMP ifIndex: 788
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, Loopback:
None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None

Logical interface xe-0/3/3.0 (Index 364) (SNMP ifIndex 829)
  Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2
  MAC address: 00:00:c8:01:01:09, Type: Configured,
    Input bytes   : 202324652
    Output bytes  : 202324560
    Input frames  : 4398362
    Output frames : 4398360
  Policer statistics:
  Policer type    Discarded frames   Discarded bytes
Output aggregate      3992386          183649756

```

Sample Output

show interfaces mc-ae

```
user@host> show interfaces mc-ae ae0 unit 512
Member Links   : ae0
Local Status   : active
Peer Status    : active
Logical Interface      : ae0.512
Core Facing Interface : Label Ethernet Interface
ICL-PL         : Label Ethernet Interface
```

show interfaces media (SONET/SDH)

The following example displays the output fields unique to the **show interfaces media** command for a SONET interface (with no level of output specified):

```
user@host> show interfaces media so-4/1/2
Physical interface: so-4/1/2, Enabled, Physical link is Up
  Interface index: 168, SNMP ifIndex: 495
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
  Loopback: None, FCS: 16, Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1783 (00:00:00 ago), Output: 1786 (00:00:08 ago)
  LCP state: Opened
  NCP state: inet: Not-configured, inet6: Not-configured, iso: Not-configured,
  mpls: Not-configured
  CHAP state: Not-configured
  CoS queues    : 8 supported
  Last flapped  : 2005-06-15 12:14:59 PDT (04:31:29 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  SONET alarms  : None
  SONET defects : None
  SONET errors:
    BIP-B1: 121, BIP-B2: 916, REI-L: 0, BIP-B3: 137, REI-P: 16747, BIP-BIP2: 0
  Received path trace: routerb so-1/1/2
  Transmitted path trace: routera so-4/1/2
```

Sample Output

show interfaces policers

```
user@host> show interfaces policers
Interface      Admin Link Proto Input Policer      Output Policer
ge-0/0/0       up    up
ge-0/0/0.0     up    up    inet
               up    up    iso
gr-0/3/0       up    up
ip-0/3/0       up    up
mt-0/3/0       up    up
pd-0/3/0       up    up
pe-0/3/0       up    up
...
so-2/0/0       up    up
so-2/0/0.0     up    up    inet  so-2/0/0.0-in-policer  so-2/0/0.0-out-policer
```

```

iso
so-2/1/0      up    down
...

```

show interfaces policers interface-name

```

user@host> show interfaces policers so-2/1/0
Interface      Admin Link Proto Input Policer      Output Policer
so-2/1/0       up    down
so-2/1/0.0     up    down inet  so-2/1/0.0-in-policer so-2/1/0.0-out-policer
                                     iso
                                     inet6

```

Sample Output

show interfaces queue

The following truncated example shows the CoS queue sizes for queues 0, 1, and 3. Queue 1 has a queue buffer size (guaranteed allocated memory) of 9192 bytes.

```

user@host> show interfaces queue
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 509
  Forwarding classes: 8 supported, 8 in use
  Egress queues: 8 supported, 8 in use
  Queue: 0, Forwarding classes: class0
    Queued:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets :                0                0 pps
      RL-dropped packets  :                0                0 pps
      RL-dropped bytes    :                0                0 bps
      RED-dropped packets :                0                0 pps
      Low               :                0                0 pps
      Medium-low         :                0                0 pps
      Medium-high        :                0                0 pps
      High               :                0                0 pps
      RED-dropped bytes   :                0                0 bps
      Low               :                0                0 bps
      Medium-low         :                0                0 bps
      Medium-high        :                0                0 bps
      High               :                0                0 bps
    Queue Buffer Usage:
      Reserved buffer    :            118750000 bytes
      Queue-depth bytes  :
      Current            :                0
    ..
    ..
  Queue: 1, Forwarding classes: class1
    ..
    ..
    Queue Buffer Usage:
      Reserved buffer    :            9192 bytes
      Queue-depth bytes  :
      Current            :                0
    ..

```



```

..
Queue: 3, Forwarding classes: class3
  Queued:
..
..
Queue Buffer Usage:
  Reserved buffer      :          6250000 bytes
  Queue-depth bytes    :
  Current              :          0
..
..

```

Sample Output

show interfaces redundancy

```

user@host> show interfaces redundancy
Interface State      Last change Primary    Secondary Current status
rsp0      Not present                sp-1/0/0  sp-0/2/0  both down
rsp1      On secondary 1d 23:56   sp-1/2/0  sp-0/3/0  primary down
rsp2      On primary   10:10:27   sp-1/3/0  sp-0/2/0  secondary down
rlsq0     On primary   00:06:24   lsq-0/3/0 lsq-1/0/0 both up

```

show interfaces redundancy (Aggregated Ethernet)

```

user@host> show interfaces redundancy
Interface State      Last change Primary    Secondary Current status
rlsq0     On secondary 00:56:12   lsq-4/0/0 lsq-3/0/0 both up

ae0
ae1
ae2
ae3
ae4

```

show interfaces redundancy detail

```

user@host> show interfaces redundancy detail
Interface      : rlsq0
State          : On primary
Last change    : 00:45:47
Primary        : lsq-0/2/0
Secondary      : lsq-1/2/0
Current status : both up
Mode           : hot-standby

Interface      : rlsq0:0
State          : On primary
Last change    : 00:45:46
Primary        : lsq-0/2/0:0
Secondary      : lsq-1/2/0:0
Current status : both up
Mode           : warm-standby

```

Sample Output

show interfaces routing brief

```

user@host> show interfaces routing brief

```

Interface	State	Addresses
so-5/0/3.0	Down	ISO enabled
so-5/0/2.0	Up	MPLS enabled
		ISO enabled
		INET 192.168.2.120
		INET enabled
so-5/0/1.0	Up	MPLS enabled
		ISO enabled
		INET 192.168.2.130
		INET enabled
at-1/0/0.3	Up	CCC enabled
at-1/0/0.2	Up	CCC enabled
at-1/0/0.0	Up	ISO enabled
		INET 192.168.90.10
		INET enabled
lo0.0	Up	ISO 47.0005.80ff.f800.0000.0108.0001.1921.6800.5061.00
		ISO enabled
		INET 127.0.0.1
fxp1.0	Up	
fxp0.0	Up	INET 192.168.6.90

show interfaces routing detail

```

user@host> show interfaces routing detail
so-5/0/3.0
  Index: 15, Refcount: 2, State: Up <Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  ISO address (null)
    State: <Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
so-5/0/2.0
  Index: 14, Refcount: 7, State: <Up Broadcast PointToPoint Multicast> Change:<>

  Metric: 0, Up/down transitions: 0, Full-duplex
  Link layer: HDLC serial line Encapsulation: PPP Bandwidth: 155Mbps
  MPLS address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4458 bytes
  ISO address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
  INET address 192.168.2.120
    State: <Up Broadcast PointToPoint Multicast Localup> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
    Local address: 192.168.2.120
    Destination: 192.168.2.110/32
  INET address (null)
    State: <Up Broadcast PointToPoint Multicast> Change: <>
    Preference: 0 (120 down), Metric: 0, MTU: 4470 bytes
...

```

Sample Output

show interfaces routing-instance all

```

user@host> show interfaces terse routing-instance all
Interface Admin Link Proto Local Remote Instance
at-0/0/1 up up inet 10.0.0.1/24

```

```

ge-0/0/0.0 up up inet 192.168.4.28/24 sample-a
at-0/1/0.0 up up inet6 fe80::a:0:0:4/64 sample-b
so-0/0/0.0 up up inet 10.0.0.1/32

```

Sample Output

show interfaces snmp-index

```

user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
Interface index: 149, SNMP ifIndex: 33
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
Link flags : Keepalives
CoS queues : 8 supported
Last flapped : 2005-06-15 11:45:57 PDT (05:38:43 ago)
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)
SONET alarms : LOL, PLL, LOS
SONET defects : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P

```

Sample Output

show interfaces source-class all

```

user@host> show interfaces source-class all
Logical interface so-0/1/0.0

      Source class          Packets          Bytes
                        (packet-per-second) (bits-per-second)
      gold                  1928095          161959980
      (                    889) (          597762)
      bronze                 0                0
      (                    0) (                0)
      silver                 0                0
      (                    0) (                0)
Logical interface so-0/1/3.0

      Source class          Packets          Bytes
                        (packet-per-second) (bits-per-second)
      gold                   0                0
      (                    0) (                0)
      bronze                 0                0
      (                    0) (                0)
      silver                116113          9753492
      (                    939) (          631616)

```

Sample Output

show interfaces statistics (Fast Ethernet)

```

user@host> show interfaces fe-1/3/1 statistics
Physical interface: fe-1/3/1, Enabled, Physical link is Up
Interface index: 144, SNMP ifIndex: 1042
Description: ford fe-1/3/1
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000

```

```

CoS queues      : 4 supported, 4 maximum usable queues
Current address: 00:90:69:93:04:dc, Hardware address: 00:90:69:93:04:dc
Last flapped   : 2006-04-18 03:08:59 PDT (00:01:24 ago)
Statistics last cleared: Never
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
Input errors: 0, Output errors: 0
Active alarms  : None
Active defects : None
Logical interface fe-1/3/1.0 (Index 69) (SNMP ifIndex 50)
  Flags: SNMP-Traps Encapsulation: ENET2
  Protocol inet, MTU: 1500
    Flags: Is-Primary, DCU, SCU-in
      Destination class      Packets      Bytes
      (packet-per-second)    (bits-per-second)
      silver1                0          0
      (                      0) (          0)
      silver2                0          0
      (                      0) (          0)
      silver3                0          0
      (                      0) (          0)
  Addresses, Flags: Is-Default Is-Preferred Is-Primary
  Destination: 10.27.245/24, Local: 10.27.245.2,
  Broadcast: 10.27.245.255
  Protocol iso, MTU: 1497
  Flags: Is-Primary

```

Sample Output

show interfaces switch-port

```

user@host# show interfaces ge-slot/0/0 switch-port port-number
Port 0, Physical link is Up
  Speed: 100mbps, Auto-negotiation: Enabled
  Statistics:
    Receive      Transmit
    Total bytes  28437086  21792250
    Total packets 409145    88008
    Unicast packets 9987      83817
    Multicast packets 145002    0
    Broadcast packets 254156    4191
    Multiple collisions 23        10
    FIFO/CRC/Align errors 0         0
    MAC pause frames 0         0
    Oversized frames 0
    Runt frames 0
    Jabber frames 0
    Fragment frames 0
    Discarded frames 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
      Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
partner Speed: 100 Mbps
    Local resolution:
      Flow control: None, Remote fault: Link OK

```

Sample Output

show interfaces transport pm

```

user@host> show interfaces transport pm all current et-0/1/0
Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current Elapse time:900 Seconds
Near End Suspect Flag:False Reason:None
PM COUNT THRESHOLD TCA-ENABLED TCA-RAISED

OTU-BBE 0 800 No No
OTU-ES 0 135 No No
OTU-SES 0 90 No No
OTU-UAS 427 90 No No
Far End Suspect Flag:True Reason:Unknown
PM COUNT THRESHOLD TCA-ENABLED TCA-RAISED

OTU-BBE 0 800 No No
OTU-ES 0 135 No No
OTU-SES 0 90 No No
OTU-UAS 0 90 No No
Near End Suspect Flag:False Reason:None
PM COUNT THRESHOLD TCA-ENABLED TCA-RAISED

ODU-BBE 0 800 No No
ODU-ES 0 135 No No
ODU-SES 0 90 No No
ODU-UAS 427 90 No No
Far End Suspect Flag:True Reason:Unknown
PM COUNT THRESHOLD TCA-ENABLED TCA-RAISED

ODU-BBE 0 800 No No
ODU-ES 0 135 No No
ODU-SES 0 90 No No
ODU-UAS 0 90 No No
FEC Suspect Flag:False Reason:None
PM COUNT THRESHOLD TCA-ENABLED TCA-RAISED

FEC-CorrectedErr 2008544300 0 NA NA
FEC-UncorrectedWords 0 0 NA NA
BER Suspect Flag:False Reason:None
PM MIN MAX AVG THRESHOLD TCA-ENABLED

TCA-RAISED
BER 3.6e-5 5.8e-5 3.6e-5 10.0e-3 No
Yes
Physical interface: et-0/1/0, SNMP ifIndex 515
14:45-current
Suspect Flag:True Reason:Object Disabled
PM CURRENT MIN MAX AVG THRESHOLD

TCA-ENABLED TCA-RAISED (MIN)

(MAX) (MIN) (MAX) (MIN) (MAX)
Lane chromatic dispersion 0 0 0 0
0 NA NA NA NA
Lane differential group delay 0 0 0 0
0 NA NA NA NA
q Value 120 120 120 120 0
0 NA NA NA NA
SNR 28 28 29 28 0
0 NA NA NA NA

```

Tx output power(0.01dBm)	-5000	-5000	-5000	-5000	-300
-100 No No No No	No				
Rx input power(0.01dBm)	-3642	-3665	-3626	-3637	-1800
-500 No No No No	No				
Module temperature(Celsius)	46	46	46	46	-5
75 No No No No	No				
Tx laser bias current(0.1mA)	0	0	0	0	0
0 NA NA NA NA	NA				
Rx laser bias current(0.1mA)	1270	1270	1270	1270	0
0 NA NA NA NA	NA				
Carrier frequency offset(MHz)	-186	-186	-186	-186	-5000
5000 No No No No	No				

Sample Output

show security zones

```

user@host> show security zones
Functional zone: management
  Description: This is the management zone.
  Policy configurable: No
  Interfaces bound: 1
  Interfaces:
    ge-0/0/0.0
Security zone: Host
  Description: This is the host zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    fxp0.0
Security zone: abc
  Description: This is the abc zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/1.0
Security zone: def
  Description: This is the def zone.
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Interfaces bound: 1
  Interfaces:
    ge-0/0/2.0

```

show interfaces snmp-index

Supported Platforms	EX Series, M Series, MX Series, PTX Series, T Series
Syntax	show interfaces snmp-index <i>snmp-index</i>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information for the interface with the specified SNMP index.
Options	This command has no options.
Additional Information	Output from both the show interfaces <i>interface-name</i> detail and the show interfaces <i>interface-name</i> extensive command includes all the information displayed in the output from the show interfaces snmp-index command.
Required Privilege Level	view
List of Sample Output	show interfaces snmp-index on page 885
Output Fields	The output fields from the show interfaces snmp-index <i>snmp-index</i> command are identical to those produced by the show interfaces <i>interface-name</i> command. For a description of output fields, see the other chapters in this manual.

Sample Output

show interfaces snmp-index

```

user@host> show interfaces snmp-index 33
Physical interface: so-2/1/1, Enabled, Physical link is Down
Interface index: 149, SNMP ifIndex: 33
Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: 0C48,
Loopback: None, FCS: 16, Payload scrambler: Enabled
Device flags   : Present Running Down
Interface flags: Hardware-Down Point-To-Point SNMP-Traps 16384
Link flags     : Keepalives
CoS queues     : 8 supported
Last flapped   : 2005-06-15 11:45:57 PDT (05:38:43 ago)
Input rate     : 0 bps (0 pps)
Output rate    : 0 bps (0 pps)
SONET alarms   : LOL, PLL, LOS
SONET defects  : LOL, PLL, LOF, LOS, SEF, AIS-L, AIS-P

```

show interfaces summary

Syntax	show interfaces summary
Release Information	Command introduced in Junos OS Release 14.1R2.
Description	Display the status and statistics on logical interfaces configured on the device at the Flexible PIC Concentrator (FPC) level.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show interfaces summary on page 886
Output Fields	Table 139 on page 886 describes the output fields for the show interfaces summary command. Output fields are listed in the approximate order in which they appear.

Table 139: show interfaces summary Output Fields

Field Name	Field Description
System's maximum logical interfaces	Total number of logical interfaces in the device.
Logical interfaces allocated	Number of allocated logical interfaces.
Logical interfaces available	Number of available logical interfaces.
Logical interface type	<p>The type of logical interfaces.</p> <ul style="list-style-type: none"> • LSI—Number of label-switched logical interfaces and their status. • Ethernet Untagged—Number of untagged logical interfaces and their status. • Ethernet VLAN—Number of tagged logical interfaces and their status. • Others—Number of dynamic and other logical interfaces, and their status.
System	Statistics on the global logical interfaces in the system.
FPC x	Statistics on the logical interfaces in a specific FPC.

Sample Output

show interfaces summary

```

user@host> show interfaces summary
Logical interfaces:
  System's maximum logical interfaces : 262144
  Logical interfaces allocated         :      31
  Logical interfaces available         : 262113

```


System:

Logical interface type	Count	UP	DOWN
Total	28	28	0
LSI	0	0	0
Ethernet Untagged	15	15	0
Ethernet VLAN	0	0	0
Others	13	13	0

FPC1:

Logical interface type	Count	UP	DOWN
Total	3	3	0
LSI	0	0	0
Ethernet Untagged	3	3	0
Ethernet VLAN	0	0	0
Others	0	0	0

FPC2:

Logical interface type	Count	UP	DOWN
Total	0	0	0
LSI	0	0	0
Ethernet Untagged	0	0	0
Ethernet VLAN	0	0	0
Others	0	0	0

show ilmi statistics

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [PTX Series](#), [T Series](#)

Syntax `show ilmi statistics`

Release Information Command introduced before Junos OS Release 7.4.

Description Display input and output Integrated Local Management Interface (ILMI) statistics.

Options This command has no options.

Required Privilege Level view

Related Documentation

- [clear ilmi statistics on page 831](#)

List of Sample Output [show ilmi statistics on page 890](#)

Output Fields [Table 140 on page 889](#) lists the output fields for the **show ilmi statistics** command. Output fields are listed in the approximate order in which they appear.

Table 140: show ilmi statistics Output Fields

Field Name	Field Description
Input	<p>Information about received ILMI packets:</p> <ul style="list-style-type: none"> • Packets—Total number of messages delivered to the ILMI entity from the transport service. • Bad versions—Total number of messages delivered to the ILMI entity that were for an unsupported ILMI version. • Bad community names—Total number of messages delivered to the ILMI entity that did not use an ILMI community name. • Bad community uses—Total number of messages delivered to the ILMI entity that represented an ILMI operation that was not allowed by the ILMI community named in the message. • ASN parse errors—Total number of ASN.1 or BER errors encountered by the ILMI entity when decoding received ILMI messages. • Too bigs—Total number of ILMI packets delivered to the ILMI entity with an error status field of tooBig. • No such names—Total number of ILMI packets delivered to the ILMI entity with an error status field of noSuchName. • Bad values—Total number of ILMI packets delivered to the ILMI entity with an error status field of badValue. • Read onlys—Total number of valid ILMI packets delivered to the ILMI entity with an error status field of readOnly. Only incorrect implementations of ILMI generate this error. • General errors—Total number of ILMI packets delivered to the ILMI entity with an error status field of genErr. • Total request varbinds—Total number of objects retrieved successfully by the ILMI entity as a result of receiving valid ILMI GetRequest and GetNext packets. • Total set varbinds—Total number of objects modified successfully by the ILMI entity as a result of receiving valid ILMI SetRequest packets. • Get requests—Total number of ILMI GetRequest packets that have been accepted and processed by the ILMI entity. • Get nexts—Total number of ILMI GetNext packets that have been accepted and processed by the ILMI entity. • Set requests—Total number of ILMI SetRequest packets that have been accepted and processed by the ILMI entity. • Get responses—Total number of ILMI GetResponse packets that have been accepted and processed by the ILMI entity. • Traps—Total number of ILMI traps received by the ILMI entity. • Silent drops—Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequest, and InformRequest packets delivered to the ILMI entity that were silently dropped because the size of a reply containing an alternate response packet with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • Proxy drops—Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequest, and InformRequest packets delivered to the ILMI entity that were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in such a way (other than a timeout) that no response packet could be returned.
Output	<p>Information about transmitted ILMI packets:</p> <ul style="list-style-type: none"> • Packets—Total number of messages passed from the ILMI entity to the transport service. • Too bigs—Total number of ILMI packets generated by the ILMI entity with an error status field of tooBig. • No such names—Total number of ILMI packets generated by the ILMI entity with an error status field of noSuchName. • Bad values—Total number of ILMI packets generated by the ILMI entity with an error status field of badValue. • General errors—Total number of ILMI packets generated by the ILMI entity with an error status field of genErr. • Get requests—Total number of ILMI GetRequest packets that have been generated by the ILMI entity. • Get nexts—Total number of ILMI GetNext packets that have been generated by the ILMI entity. • Set requests—Total number of ILMI SetRequest packets that have been generated by the ILMI entity. • Get responses—Total number of ILMI GetResponse packets that have been generated by the ILMI entity. • Traps—Total number of ILMI traps generated by the ILMI entity.

Sample Output

show ilmi statistics

```
user@host> show ilmi statistics
ILMI statistics:
Input:
  Packets: 0, Bad versions: 0, Bad community names: 0,
  Bad community uses: 0, ASN parse errors: 0,
  Too bigs: 0, No such names: 0, Bad values: 0,
  Read onlys: 0, General errors: 0,
  Total request varbinds: 0, Total set varbinds: 0,
  Get requests: 0, Get nexts: 0, Set requests: 0,
  Get responses: 0, Traps: 0,
  Silent drops: 0, Proxy drops 0
Output:
  Packets: 0, Too bigs: 0, No such names: 0,
  Bad values: 0, General errors: 0,
  Get requests: 0, Get nexts: 0, Set requests: 0,
  Get responses: 0, Traps: 0
```

show security alarms

Supported Platforms [SRX Series](#), [vSRX](#)

Syntax show security alarms
 <detail>
 <alarm-id *id-number*>
 <alarm-type [*types*]>
 <newer-than YYYY-MM-DD.HH:MM:SS>
 <older-than YYYY-MM-DD.HH:MM:SS>
 <process *process*>
 <severity *severity*>

Release Information Command introduced in Junos OS Release 11.2.

Description Display the alarms that are active on the device. Run this command when the CLI prompt indicates that a security alarm has been raised, as shown here:

```
[1 SECURITY ALARM] user@host#
```

Options **none**—Display all active alarms.

detail—(Optional) Display detailed output.

alarm-id *id-number*—(Optional) Display the specified alarm.

alarm-type [*types*]—(Optional) Display the specified alarm type or a set of types.

You can specify one or more of the following alarm types:

- authentication
- cryptographic-self-test
- decryption-failures
- encryption-failures
- ike-phase1-failures
- ike-phase2-failures
- key-generation-self-test
- non-cryptographic-self-test
- policy
- replay-attacks

newer-than YYYY-MM-DD.HH:MM:SS—(Optional) Display active alarms that were raised after the specified date and time.

older-than YYYY-MM-DD.HH:MM:SS—(Optional) Display active alarms that were raised before the specified date and time.

process *process*—(Optional) Display active alarms that were raised by the specified system process.

severity *severity*—(Optional) Display active alarms of the specified severity.

You can specify the following severity levels:

- **alert**
- **crit**
- **debug**
- **emerg**
- **err**
- **info**
- **notice**
- **warning**

Required Privilege Level security—To view this statement in the configuration.

Related Documentation

- [clear security alarms](#)
- [Example: Generating a Security Alarm in Response to Policy Violations](#)

List of Sample Output

[show security alarms on page 893](#)
[show security alarms detail on page 893](#)
[show security alarms alarm-id on page 893](#)
[show security alarms alarm-type authentication on page 894](#)
[show security alarms newer-than <time> on page 894](#)
[show security alarms older-than <time> on page 894](#)
[show security alarms process <process> on page 894](#)
[show security alarms severity <severity> on page 894](#)

Output Fields [Table 141 on page 892](#) lists the output fields for the **show security alarms** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used.

Table 141: show security alarms

Field Name	Field Description	Level of Output
ID	Identification number of the alarm.	All levels
Alarm time	Date and time the alarm was raised..	All levels

Table 141: show security alarms (*continued*)

Field Name	Field Description	Level of Output
Message	Information about the alarm, including the alarm type, username, IP address, and port number.	All levels
Process	System process (For example, login or sshd) and process identification number associated with the alarm.	detail
Severity	Severity level of the alarm.	detail

Sample Output

show security alarms

```
[3 SECURITY ALARMS] user@router> show security alarms
```

```

ID      Alarm time      Message
1       2010-01-19 13:41:36 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
failures (1) for user 'user' reached from '203.0.113.2'
2       2010-01-19 13:41:52 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
failures (1) for user 'user' reached from '203.0.113.2'
3       2010-01-19 13:42:13 PST  SSHD_LOGIN_FAILED_LIMIT: Specified number of login
failures (1) for user 'user' reached from '203.0.113.2'
```

show security alarms detail

```
[3 SECURITY ALARMS] user@router> show security alarms detail
```

```

Alarm ID   : 1
Alarm Type : authentication
Time       : 2010-01-19 13:41:36 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for
user 'user' reached from '203.0.113.2'
Process    : sshd (pid 1414)
Severity   : notice

Alarm ID   : 2
Alarm Type : authentication
Time       : 2010-01-19 13:41:52 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for
user 'user' reached from '203.0.113.2'
Process    : sshd (pid 1414)
Severity   : notice

Alarm ID   : 3
Alarm Type : authentication
Time       : 2010-01-19 13:42:13 PST
Message    : SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for
user 'user' reached from '203.0.113.2'
Process    : sshd (pid 1414)
Severity   : notice
```

show security alarms alarm-id

```
[3 SECURITY ALARMS] user@router> show security alarms alarm-id 1
```

ID	Alarm time	Message
1	2010-01-19 13:41:36 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'

show security alarms alarm-type authentication

```
[3 SECURITY ALARMS] user@router> show security alarms alarm-type authentication
```

ID	Alarm time	Message
1	2010-01-19 13:41:36 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
2	2010-01-19 13:41:52 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
3	2010-01-19 13:42:13 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'

show security alarms newer-than <time>

```
[3 SECURITY ALARMS] user@router> show security alarms newer-than 2010-01-19.13:41:59
```

3	2010-01-19 13:42:13 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
---	-------------------------	--

show security alarms older-than <time>

```
[3 SECURITY ALARMS] user@router> show security alarms older-than 2010-01-19.13:41:59
```

ID	Alarm time	Message
1	2010-01-19 13:41:36 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
2	2010-01-19 13:41:52 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'

show security alarms process <process>

```
[3 SECURITY ALARMS] user@router> show security alarms process sshd
```

ID	Alarm time	Message
1	2010-01-19 13:41:36 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
2	2010-01-19 13:41:52 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
3	2010-01-19 13:42:13 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'

show security alarms severity <severity>

```
[3 SECURITY ALARMS] user@router> show security alarms severity notice
```

ID	Alarm time	Message
1	2010-01-19 13:41:36 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
2	2010-01-19 13:41:52 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'
3	2010-01-19 13:42:13 PST	SSHD_LOGIN_FAILED_LIMIT: Specified number of login failures (1) for user 'user' reached from '203.0.113.2'

show security datapath-debug capture

Supported Platforms [SRX5400, SRX5600, SRX5800](#)

Syntax `show security datapath-debug capture`

Release Information Command introduced in Junos OS Release 10.0.

Description Display details of the data path debugging capture file.

Required Privilege Level view

Related Documentation

- [show security datapath-debug counter on page 896](#)
- *Understanding Data Path Debugging for Logical Systems*

List of Sample Output [show security datapath—debug capture on page 895](#)

Output Fields Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath—debug capture

```
user@host> show security datapath-debug capture
Packet 1, len 120: (C0/F0/P0/SEQ:71:1bt)
91 00 00 47 11 00 10 00 9a 14 00 19 03 00 00 00
00 00 00 00 00 01 00 47 10 00 00 00 00 00 00 00
00 1f 12 f8 dd 29 00 21 59 84 f4 01 81 00 02 1e
08 00 45 60 01 f4 00 00 00 00 3f 06 73 9f 01 01
01 02 03 01 01 02 d4 31 d4 31 00 00 00 00 00 00
00 00 50 02 00 00 ff ad 00 00 00 00
Packet 2, len 120: (C0/F0/P0/SEQ:71:1bt)
90 00 00 47 04 00 00 00 00 00 00 00 02 02 00 47
10 00 00 00 00 00 00 00 50 00 a6 1c 00 00 00 00
00 00 00 0a 00 00 00 00 00 00 09 d9 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 1f 12 f8
dd 29 00 21 59 84 f4 01 81 00 02 1e
```

show security datapath-debug counter

Supported Platforms [SRX5400, SRX5600, SRX5800](#)

Syntax show security datapath-debug counter

Release Information Command introduced in Junos OS Release 10.0.

Description Display details of the data path debugging counter.

Required Privilege Level view

Related Documentation

- [show security datapath-debug capture on page 895](#)
- *Understanding Data Path Debugging for Logical Systems*

List of Sample Output [show security datapath-debug counter on page 896](#)

Output Fields Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath-debug counter

```
user@host> show security datapath-debug counter
Datapath debug counters
Packet Filter 1:
np-ingress
Chassis 0 FPC 4 : 1
np-ingress
Chassis 0 FPC 3 : 0
np-egress
Chassis 0 FPC 4 : 1
np-egress
Chassis 0 FPC 3 : 0
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 2
pot
Chassis 0 FPC 0 PIC 1: 0
pot
```

show security monitoring

Supported Platforms [SRX Series, vSRX](#)

Syntax show security monitoring

Release Information Command introduced in Junos OS Release 10.2.

Description Displays a count of security flow and central point (CP) sessions, CPU utilization (as a percentage of maximum), and memory in use (also as a percentage of maximum) at the moment the command is run. This command is supported on SRX1500, SRX5400, SRX5600, and SRX5800 devices and vSRX.

Required Privilege Level View

Related Documentation

- [show security monitoring fpc fpc-number on page 899](#)
- [show security monitoring performance session on page 902](#)
- [show security monitoring performance spu on page 903](#)

show security monitoring

```
user@host>show security monitoring
```

```
user@host> show security monitoring
```

FPC	PIC	CPU	Mem	Flow session current	Flow session maximum	CP session current	CP session maximum
1	0	0	11	0	0	0	0
1	1	0	5	3	6291456	1	7549747
1	2	0	5	2	6291456	0	7549747
1	3	0	5	3	6291456	1	7549747
8	0	0	65	4	6963	2	8355
8	1	0	65	2	6963	0	8355
Total Sessions:				14	18888294	4	22665951

show security monitoring (vSRX)

```
user@host>show security monitoring
```

```
user@host> show security monitoring
```

FPC	PIC	CPU	Mem	Flow session current	Flow session maximum	CP session current	CP session maximum
0	0	0	68	2	524288	N/A	N/A

show security monitoring (vSRX in a Chassis Cluster)

```
user@host>show security monitoring
```

```
node0:
```

FPC	PIC	CPU	Mem	Flow session current	Flow session maximum	CP session current	CP session maximum
0	0	0	67	0	524288	N/A	N/A

```
node1:
```

FPC	PIC	CPU	Mem	Flow session current	Flow session maximum	CP session current	CP session maximum
0	0	0	67	0	524288	N/A	N/A

show security monitoring fpc fpc-number

Supported Platforms [SRX Series, vSRX](#)

Syntax `show security monitoring fpc fpc-number`
`<node (node-id | all | local | primary)>`

Release Information Command introduced in Junos OS Release 9.2.

Description Display security monitoring information about the FPC slot.

- Options**
- **fpc-number**—Display security monitoring information for the specified FPC slot. It can be in the range from 0 to 11.
 - **node**—(Optional) For chassis cluster configurations, display security monitoring information for the specified FPC on a specific node (device) in the cluster.
 - **node-id**—Identification number of the node. It can be 0 or 1.
 - **all**—Display information about all nodes.
 - **local**—Display information about the local node.
 - **primary**—Display information about the primary node.

Additional Information For complete list of slot numbering, physical port, and logical interface numbering for SRX Series devices in chassis cluster, see *Chassis Cluster Feature Guide for SRX Series Devices*.

Required Privilege Level view

Related Documentation • [show services ip-monitoring status on page 905](#)

List of Sample Output [show security monitoring fpc 0 on page 900](#)
[show security monitoring fpc 1 on page 900](#)
[show security monitoring fpc 8 on page 901](#)

Output Fields [Table 142 on page 899](#) lists the output fields for the `show security monitoring fpc fpc-number` command. Output fields are listed in the approximate order in which they appear.

Table 142: show security monitoring fpc fpc-number Output Fields

Field Name	Field Description
FPC	Slot number in which the FPC is installed.
PIC	Slot number in which the PIC is installed.

Table 142: show security monitoring fpc fpc-number Output Fields (*continued*)

Field Name	Field Description
CPU Utilization (%)	Total percentage of CPU being used by the PIC's processors.
Memory Utilization (%)	Percentage of heap space (dynamic memory) being used by the PIC's processor. If this number exceeds 80 percent, there might be a software problem (memory leak).
Current flow session	The current number of flow sessions. When SRX Series devices operate in packet mode, flow sessions will not be created and this field will remain zero.
Max flow session	The maximum number of flow sessions allowed. This number will differ from one device to another.
SPU current cp session	The current number of cp sessions for the SPU (on SRX5600, and SRX5800 devices only).
SPU max cp session	The maximum number of cp sessions allowed for the SPU (on SRX5600, and SRX5800 devices only).

Sample Output

show security monitoring fpc 0

```

user@host> show security monitoring fpc 0
FPC 0
  PIC 0
    CPU utilization      :    0 %
    Memory utilization   :   82 %
    Current flow session :    0
    Max flow session     :    0
    Current CP session   :    0
    Max CP session       : 12000000
  Session Creation Per Second (for last 96 seconds on average):    0
  PIC 1
    CPU utilization      :    0 %
    Memory utilization   :   54 %
    Current flow session :    0
    Max flow session     : 819200
    Current CP session   :    0
    Max CP session       :    0
  Session Creation Per Second (for last 96 seconds on average):    0

```

Sample Output

show security monitoring fpc 1

```

user@host> show security monitoring fpc 1
FPC 1
  PIC 0
    CPU utilization      :    0 %
    Memory utilization   :   21 %
    Current flow session :    0
    Max flow session     : 524288
    Current CP session   :    0

```

```
Max CP session      : 1048576
Session Creation Per Second (for last 96 seconds on average): 0
```

Sample Output

show security monitoring fpc 8

```
user@host> show security monitoring fpc 5
FPC 5
  PIC 0
    CPU utilization      : 0 %
    Memory utilization   : 64 %
    Current flow session : 0
    Max flow session     : 524288
    Current CP session   : 0
    Max CP session       : 2359296
  Session Creation Per Second (for last 96 seconds on average): 0
  PIC 1
    CPU utilization      : 0 %
    Memory utilization   : 65 %
    Current flow session : 0
    Max flow session     : 1048576
    Current CP session   : 0
    Max CP session       : 0
  Session Creation Per Second (for last 96 seconds on average): 0
```

show security monitoring performance session

Supported Platforms [SRX Series, vSRX](#)

Syntax show security monitoring performance session

<fpc slot-number>

<pic slot-number>

Release Information Command introduced in Junos OS Release of 10.2.

Description Display the current session (total number of sessions at that time) for the last 60 seconds.

- Options**
- **fpc slot-number** — Display information about the FPC slot. Use this option to filter the output based on the slot number.
 - **pic slot-number** — Display information about existing PIMs or Mini-PIMs in a particular PIC slot. Use this option to filter the output based on PIC slot.



NOTE: The fpc slot-number and pic slot-number options are not available on SRX300, SRX320, and SRX340 devices.

Required Privilege Level View

Related Documentation

- [show services ip-monitoring status on page 905](#)

show security monitoring performance session

user@host> show security monitoring performance session

```
fpc 0 pic 0
Last 60 seconds:
0:      8  1:      8  2:      8  3:      8  4:      8  5:      7
6:      7  7:      7  8:      7  9:      7 10:      7 11:      8
12:      8 13:      8 14:      7 15:      7 16:      7 17:      7
18:      7 19:      7 20:      7 21:      5 22:      5 23:      5
24:      5 25:      5 26:      5 27:      5 28:      5 29:      4
30:      4 31:      4 32:      3 33:      3 34:      3 35:      3
36:      5 37:      5 38:      6 39:      6 40:      5 41:      5
42:      5 43:      5 44:      5 45:      5 46:      5 47:      5
48:      7 49:      7 50:      6 51:      8 52:      8 53:      6
54:      5 55:      7 56:      7 57:      5 58:      5 59:      8
```


show security monitoring performance spu

Supported Platforms [SRX Series, vSRX](#)

Syntax show security monitoring performance spu

<fpc slot-number>

<pic slot-number>

Release Information Command introduced in Junos OS Release 10.2.

Description Display the services processing unit (SPU) percent utilization for all FPC slots over the last 60 seconds. Use this command to track the percent utilization statistics per second for the past 60 seconds for each FPC slot and PIC.

- Options**
- **fpc slot-number** — Display information about the FPC slot. Use this option to filter the output based on the slot number.
 - **pic slot-number** — Display information about existing PIMs or Mini-PIMs in a particular PIC slot. Use this option to filter the output based on PIC slot.



NOTE: The fpc slot-number and pic slot-number options are not available on SRX300, SRX320, or SRX340 devices or on vSRX instances.

Required Privilege Level View

Related Documentation

- [show services ip-monitoring status on page 905](#)

show security monitoring performance spu

This sample shows 46% utilization of the SPU for second 42 in the past 60 seconds for FPC 0 and PIC 0.

```
user@host>show security monitoring performance spu
```

```
fpc 0 pic 0
Last 60 seconds:
 0: 48  1: 48  2: 48  3: 48  4: 48  5: 48
 6: 48  7: 48  8: 49  9: 48 10: 48 11: 48
12: 48 13: 48 14: 48 15: 48 16: 48 17: 48
18: 48 19: 48 20: 48 21: 48 22: 49 23: 48
24: 49 25: 49 26: 48 27: 48 28: 48 29: 48
30: 48 31: 48 32: 48 33: 48 34: 48 35: 48
```

36:	46	37:	47	38:	46	39:	46	40:	46	41:	46
42:	46	43:	46	44:	46	45:	46	46:	46	47:	46
48:	46	49:	46	50:	46	51:	46	52:	46	53:	46
54:	46	55:	46	56:	46	57:	46	58:	46	59:	46

show services ip-monitoring status

Supported Platforms [SRX Series, vSRX](#)

Syntax `show services ip-monitoring status`

Release Information Command modified in Junos OS Release 11.4 R2. Next-hop functionality added in Junos OS Release 12.1X46-D15.

Description Display a brief summary of IP monitoring status along with the current state for a given policy.

Required Privilege Level view

Related Documentation

- [show security monitoring on page 897](#)

List of Sample Output

- [show services ip-monitoring status on page 906](#)
- [show services ip-monitoring status on page 906](#)
- [show services ip-monitoring status on page 907](#)
- [show services ip-monitoring status on page 907](#)
- [show services ip-monitoring status on page 907](#)

Output Fields [Table 143 on page 905](#) lists the output fields for the **show services ip-monitoring status** command. Output fields are listed in the approximate order in which they appear.

Table 143: show services ip-monitoring status Output Fields

Field Name	Field Description
Policy	Name of the policy configured.
Probe Name	Name of the probe configured.
Address	Displays the configured target address.
Status	Displays the status of the probe on the target address. If the status is PASS, then the target address is reached.
Route-Action	Displays route injection information configured for the policy and its failover status.
Route-Instance	Displays the routing instance of the route to be injected during failover.
Route	Routing address of the route to be injected during failover.
Next-Hop	Specifies the next-hop address of the route to be injected during failover. P2P interfaces only.

Table 143: show services ip-monitoring status Output Fields (*continued*)

Field Name	Field Description
State	Display the state of the route injection action. If the state is APPLIED, then the ip-monitoring policy is in failover state.
Interface Action	Displays the interface action type as enable or disable.
Policy Action	Displays the policy action type as enable or disable.
Admin State	Displays the current admin state of the interface.
Action Status	Displays the current action status of the interface.

Sample Output

show services ip-monitoring status

```

user@host> show services ip-monitoring status

Policy - policy1 (Non-preemptive. Status: FAIL)
RPM Probes:
  Probe name      Test Name      Address      Status
  -----
  probe_a         a1             15.1.1.10    FAIL
  probe_a         a2             200.1.1.1    FAIL
Route-Action:
  route-instance  route          next-hop      State
  -----
  inet.0          200.1.1.0      150.1.1.1    APPLIED
Interface-Action:
  interface      policy action  admin state  action status
  -----
  fe-0/0/5.2     Enable        UP           FAILOVER
  fe-0/0/5.4     Disable       DOWN         FAILOVER
  t1-1/0/0        Enable        UP           FAILOVER
  d10             Enable        UP           FAILOVER
  ge-0/0/1        Enable        UP           FAILOVER

```

Sample Output

show services ip-monitoring status

In this example, the policy is in the failback state, and the no-preempt option is not configured.

```

user@host> show services ip-monitoring status

Policy - policy1 (Status: PASS)
RPM Probes:
  Probe name      Test Name      Address      Status
  -----
  probe1          a1             99.1.1.2     PASS

```

```

Route-Action:
route-instance  route      next-hop      state
-----
inet.0          99.1.1.0      12.12.12.2  NOT-APPLIED
Interface-Action:
interface      policy action  admin state  action status
-----
at-2/0/0      Enable        DOWN        MARKED-DOWN
ge-0/0/2.2    Enable        DOWN        MARKED-DOWN
ge-0/0/2.3    Enable        DOWN        MARKED-DOWN

```

Sample Output

show services ip-monitoring status

In this example, the policy is in the failover state, and the primary is restored. The no-preempt option is configured.

```
user@host> show services ip-monitoring status
```

```

Policy - policy1 (Non-preemptive. Status: FAILOVER-NO-PREEMPT)
RPM Probes:
  Probe name      Test Name      Address      Status
  -----
  probe1          a1             99.1.1.2     PASS
Route-Action:
route-instance  route      next-hop      state
-----
inet.0          99.1.1.0      12.12.12.2  APPLIED
Interface-Action:
interface      policy action  admin state  action status
-----
at-2/0/0      Enable        UP          FAILOVER
ge-0/0/2.2    Enable        UP          FAILOVER
ge-0/0/2.3    Enable        UP          FAILOVER

```

Sample Output

show services ip-monitoring status

When the probe succeeds and the policy is not applied, the output is as follows:

```

user@host> show services ip-monitoring status
Policy payment (Status: PASS)
RPM Probes:
  Probe name      Test Name      Address      Status
  -----
  Probe-Payment-Server  paysvr      9.9.9.2     PASS
Route-Action:
route-instance  route      next-hop      state
-----
inet.0          9.9.9.0/24  e1-6/0/0.0    NOT-APPLIED

```

Sample Output

show services ip-monitoring status

When the probe fails and the policy is applied, the output is as follows:

```
user@host> show services ip-monitoring status
```

```
Policy payment (Status: FAIL)
```

```
RPM Probes:
```

Probe name	Test Name	Address	Status
Probe-Payment-Server	paysvr	9.9.9.2	FAIL

```
Route-Action:
```

route-instance	route	next-hop	state
inet.0	9.9.9.0/24	e1-6/0/0.0	APPLIED

show services rpm twamp client connection

Supported Platforms [SRX Series](#)

Syntax `show services rpm twamp client connection`
`<connection-name>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify a control connection name when you issue the command.

Options *connection-name*—(Optional) Display information about the specified control connection or TWAMP control client.

Required Privilege Level view

List of Sample Output [show services rpm twamp client connection on page 910](#)

Output Fields [Table 144 on page 909](#) lists the output fields for the `show services rpm twamp client connection` command. Output fields are listed in the approximate order in which they appear.

Table 144: show services rpm twamp client connection Output Fields

Field Name	Field Description
Connection Name	Connection name that uniquely identifies the connection between the TWAMP server and a particular client.
Client address	Client IP address.
Client port	Client port number.
Server address	Server IP address.
Server port	Server port number.
Session count	Session count.
Auth mode	Authentication mode.

Sample Output

show services rpm twamp client connection

```
user@host> show services rpm twamp client connection
Connection      Client      Client Server      Server Session Auth
Name            address    port  address    port  count  mode
c1              123.0.0.2  60530 123.0.0.1  1862    1
Unauthenticated
```


show services rpm twamp client history-results

Supported Platforms [SRX Series](#)

Syntax show services rpm history-results
 <brief | detail>
 <control-connection *control-connection-name*>
 <since *time*>
 <test-session *test-session-name*>

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) instance. You can also view the historical results of the probes or test packets sent from a TWAMP client to a TWAMP server for a particular control connection, or a test session associated with a control connection.

Options **brief | detail**—(Optional) Display the specified level of output.

control-connection *control-connection-name*—(Optional) Display information for the specified control connection between a TWAMP client and a TWAMP server.

since *time*—(Optional) Display information from the specified time. Specify time as *yyyy-mm-dd.hh:mm:ss*.

test-session *test-session-name*—(Optional) Display information for the specified test session associated with a control connection between a TWAMP client and a TWAMP server.

Required Privilege Level view

List of Sample Output [show services rpm twamp client history-results on page 912](#)

Output Fields [Table 145 on page 911](#) lists the output fields for the **show services rpm twamp client history-results** command. Output fields are listed in the approximate order in which they appear.

Table 145: show services rpm twamp client history-results Output Fields

Field Name	Field Description
Owner	Probe owner or the TWAMP client.
Test	Name of a test for a TWAMP probe instance.
Probe received	Timestamp when the probe result was received.

Table 145: show services rpm twamp client history-results Output Fields (*continued*)

Field Name	Field Description
Round trip time	Average ping round-trip time (RTT), in microseconds.
Probes sent	Timestamp when the probe result was sent.
Round trip time	Average ping RTT, in microseconds.

Sample Output

show services rpm twamp client history-results

```

user@host> show services rpm twamp client history-results
Owner, Test          Probe Sent          Probe received
Round trip time
c2, t2              Fri Jul 21 05:11:06 2017 Fri Jul 21 05:11:06 2017
                241 usec
c2, t2              Fri Jul 21 05:11:07 2017 Fri Jul 21 05:11:07 2017
                254 usec
c2, t2              Fri Jul 21 05:11:08 2017 Fri Jul 21 05:11:08 2017
                248 usec
c2, t2              Fri Jul 21 05:11:09 2017 Fri Jul 21 05:11:09 2017
                241 usec
c2, t2              Fri Jul 21 05:11:10 2017 Fri Jul 21 05:11:10 2017
                245 usec
c2, t2              Fri Jul 21 05:11:11 2017 Fri Jul 21 05:11:11 2017
                231 usec
c2, t2              Fri Jul 21 05:11:12 2017 Fri Jul 21 05:11:12 2017
                258 usec
c2, t2              Fri Jul 21 05:11:13 2017 Fri Jul 21 05:11:13 2017
                232 usec
c2, t2              Fri Jul 21 05:11:14 2017 Fri Jul 21 05:11:14 2017
                236 usec
c2, t2              Fri Jul 21 05:11:15 2017 Fri Jul 21 05:11:15 2017
                236 usec
c2, t2              Fri Jul 21 05:11:16 2017 Fri Jul 21 05:11:16 2017
                213 usec

```

show services rpm twamp client probe-results

Supported Platforms [SRX Series](#)

Syntax `show services rpm twamp client probe-results`
`<control-connection control-connection-name>`
`<test-session test-session-name>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Display the results of the most recent real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) probes sent from the TWAMP client to the TWAMP server. You can also view the results of the probes or test packets sent from a TWAMP client to a TWAMP server for a particular control connection, or a test session associated with a control connection.

Options **control-connection *control-connection-name***—(Optional) Display information for the specified control connection between a TWAMP client and a TWAMP server.

test-session *test-session-name*—(Optional) Display information for the specified test session associated with a control connection between a TWAMP client and a TWAMP server.

Required Privilege Level view

List of Sample Output [show services rpm twamp client probe-results on page 916](#)

Output Fields [Table 146 on page 913](#) lists the output fields for the **show services twamp client probe-results** command. Output fields are listed in the approximate order in which they appear.

Table 146: show services twamp client probe-results Output Fields

Field Name	Field Description
Owner	Name of the session-sender or the control client, which is the TWAMP client. When you configure the control-client-name option at the [edit services twamp client control-connection] hierarchy level, this field displays the configured owner name or the client name.
Test	Name of a test representing a collection of probes. When you configure the test-session test-name statement at the [edit services owner] hierarchy level, the field displays the configured test name.
server-address	Destination address used for the probes.
server-port	Destination port used for the probes.
Client address	Source or TWAMP client address used for the probes.

Table 146: show services twamp client probe-results Output Fields (*continued*)

Field Name	Field Description
Client port	Source or TWAMP client port used for the probes.
Reflector address	Session reflector or TWAMP server address used for the probes.
Reflector port	Session reflector or TWAMP server port used for the probes.
Test size	Number of probes within a test.
Probe results	<p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> • Response received—Timestamp when the probe result was determined. • Client and server hardware timestamps—If timestamps are configured, an entry appears at this point. • Rtt—Average ping round-trip time (RTT), in microseconds. • Egress jitter—Egress jitter, in microseconds. • Ingress jitter—Ingress jitter, in microseconds. • Round trip jitter—Round-trip jitter, in microseconds. • Egress interarrival jitter—Egress interarrival jitter, in microseconds. • Ingress interarrival jitter—Ingress interarrival jitter, in microseconds. • Round trip interarrival jitter—Round-trip interarrival jitter, in microseconds.
Results over current test	<p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent within the current test. • Probes received—Number of probe responses received within the current test. • Loss percentage—Percentage of lost probes for the current test. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum.

Table 146: show services twamp client probe-results Output Fields (*continued*)

Field Name	Field Description
Results over last test	<p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed:</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent for the most recently completed test. • Probes received—Number of probe responses received for the most recently completed test. • Loss percentage—Percentage of lost probes for the most recently completed test. • Test completed—Time the most recent test was completed. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individually calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test. • Maximum—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test. • Average—Average RTT, ingress delay, or egress delay measured for the most recently completed test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum.
Results over all tests	<p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent in all tests. • Probes received—Number of probe responses received in all tests. • Loss percentage—Percentage of lost probes in all tests. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types icmp-ping-timestamp and udp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individually calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum.

Sample Output

show services rpm twamp client probe-results

```
user@host> show services rpm twamp client probe-results
Owner: c2, Test: t2
  server-address: 23.0.0.1, server-port: 862, Client address: 23.0.0.1,
  Client port: 63223
  Reflector address: 23.0.0.1, Reflector port: 10000,
  Sender address: 23.0.0.1, sender-port: 10000
  Test size: 400000 probes
  Probe results:
    Response received
    Fri Jul 21 05:12:12 2017
    Fri Jul 21 05:12:12 2017
    Rtt: 242 usec, Egress jitter: -9 usec, Ingress jitter: 24 usec,
    Round trip jitter: 15 usec, Egress interarrival jitter: 14 usec,
    Ingress interarrival jitter: 18 usec,
    Round trip interarrival jitter: 16 usec
  Results over current test:
    Probes sent: 67, Probes received: 67, Loss percentage: 0.000000
    Measurement: Round trip time
  Samples: 67, Minimum: 206 usec, Maximum: 266 usec, Average: 243 usec,
    Peak to peak: 60 usec, Stddev: 11 usec, Sum: 16283 usec
    Measurement: Negative egress jitter
    Samples: 66, Minimum: 4 usec, Maximum: 25 usec, Average: 14 usec,
    Peak to peak: 21 usec, Stddev: 4 usec, Sum: 935 usec
    Measurement: Positive ingress jitter
    Samples: 55, Minimum: 0 usec, Maximum: 62 usec, Average: 19 usec,
    Peak to peak: 62 usec, Stddev: 13 usec, Sum: 1023 usec
    Measurement: Negative ingress jitter
    Samples: 11, Minimum: 1 usec, Maximum: 23 usec, Average: 8 usec,
    Peak to peak: 22 usec, Stddev: 6 usec, Sum: 87 usec
    Measurement: Positive round trip jitter
    Samples: 33, Minimum: 0 usec, Maximum: 49 usec, Average: 14 usec,
    Peak to peak: 49 usec, Stddev: 11 usec, Sum: 463 usec
    Measurement: Negative round trip jitter
    Samples: 33, Minimum: 1 usec, Maximum: 34 usec, Average: 14 usec,
    Peak to peak: 33 usec, Stddev: 8 usec, Sum: 462 usec
  Results over all tests:
    Probes sent: 67, Probes received: 67, Loss percentage: 0.000000
    Measurement: Round trip time
    Samples: 67, Minimum: 206 usec, Maximum: 266 usec, Average: 243 usec,
    Peak to peak: 60 usec, Stddev: 11 usec, Sum: 16283 usec
    Measurement: Negative egress jitter
    Samples: 66, Minimum: 4 usec, Maximum: 25 usec, Average: 14 usec,
    Peak to peak: 21 usec, Stddev: 4 usec, Sum: 935 usec
    Measurement: Positive ingress jitter
    Samples: 55, Minimum: 0 usec, Maximum: 62 usec, Average: 19 usec,
    Peak to peak: 62 usec, Stddev: 13 usec, Sum: 1023 usec
    Measurement: Negative ingress jitter
    Samples: 11, Minimum: 1 usec, Maximum: 23 usec, Average: 8 usec,
    Peak to peak: 22 usec, Stddev: 6 usec, Sum: 87 usec
    Measurement: Positive round trip jitter
    Samples: 33, Minimum: 0 usec, Maximum: 49 usec, Average: 14 usec,
    Peak to peak: 49 usec, Stddev: 11 usec, Sum: 463 usec
    Measurement: Negative round trip jitter
    Samples: 33, Minimum: 1 usec, Maximum: 34 usec, Average: 14 usec,
    Peak to peak: 33 usec, Stddev: 8 usec, Sum: 462 usec
```


show services rpm twamp client session

Supported Platforms [SRX Series](#)

Syntax `show services rpm twamp client session`
`<control-connection control-connection-name>`
`<test-session test-session-name>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients for control packets and data packets. By default, all established control connection and data connection or test sessions are displayed, unless you specify a control connection name or a test session name when you issue the command.

Options `control-connection control-connection-name`—(Optional) Display information about the specified control connection, which is established for control packets exchanged between a TWAMP client and a TWAMP server.

`test-session test-session-name`—(Optional) Display information about the specified test session, which is established for data packets transmitted between a TWAMP client and a TWAMP server, associated with a control connection.

Required Privilege Level view

List of Sample Output [show services rpm twamp client session on page 919](#)

Output Fields [Table 147 on page 918](#) lists the output fields for the `show services rpm twamp client session` command. Output fields are listed in the approximate order in which they appear.

Table 147: show services rpm twamp client session Output Fields

Field Name	Field Description
Connection Name	Name of the control connection that uniquely identifies the connection between the TWAMP server and the TWAMP client.
Session Name	Name of the test session that uniquely identifies the data session between the TWAMP server and the TWAMP client.
Sender address	Sender IP address.
Sender port	Sender port number.
Reflector address	Reflector IP address.
Reflector port	Reflector port number.

Sample Output

show services rpm twamp client session

```
user@host> show services rpm twamp client session
root> show services rpm twamp client session
```

Connection Name	Session Name	Sender address	Sender port	Reflector address	Reflector port
c1	t1	123.0.0.2	10001	123.0.0.1	10001

show services rpm twamp server connection

Supported Platforms [SRX Series](#)

Syntax `show services rpm twamp server connection
<connection-id>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command.

Options *connection-id*—(Optional) Identifier of the connection that you want to display information about.

Required Privilege Level view

List of Sample Output [show services rpm twamp server connection on page 921](#)

Output Fields [Table 144 on page 909](#) lists the output fields for the `show services rpm twamp server connection` command. Output fields are listed in the approximate order in which they appear.

Table 148: show services rpm twamp server connection Output Fields

Field Name	Field Description
Connection ID	Connection ID that uniquely identifies the connection between the TWAMP server and a particular client.
Client address	Client IP address.
Client port	Client port number.
Server address	Server IP address.
Server port	Server port number.
Session count	Session count.
Auth mode	Authentication mode.

Sample Output

show services rpm twamp server connection

```
user@host> show services rpm twamp server connection
Connection Client      Client  Server  Server  Session  Auth
ID          address      port    address port    count    mode
          5 123.0.0.1      50821 123.0.0.2 862      1
Unauthenticated
```

show services rpm twamp server session

Supported Platforms [SRX Series](#)

Syntax `show services rpm twamp server session`
`<session-id>`

Release Information Command introduced in Junos OS Release 15.1X49-D110 for SRX Series devices.

Description Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command.

Options *session-id*—(Optional) Identifier of the session that you want to display information about.

Required Privilege Level view

List of Sample Output [show services rpm twamp server session on page 922](#)

Output Fields [Table 149 on page 922](#) lists the output fields for the `show services rpm twamp server session` command. Output fields are listed in the approximate order in which they appear.

Table 149: show services rpm twamp server session Output Fields

Field Name	Field Description
Session ID	Session ID that uniquely identifies the session between the TWAMP server and a particular client.
Connection ID	Connection ID that uniquely identifies the connection between the TWAMP server and a particular client.
Sender address	Sender IP address.
Sender port	Sender port number.
Reflector address	Reflector IP address.
Reflector port	Reflector port number.

Sample Output

`show services rpm twamp server session`

```
user@host> show services rpm twamp server session
```

Session ID	Connection ID	Sender address	Sender port	Reflector address	Reflector port
3	6	123.0.0.1	40000	123.0.0.2	40000

show snmp health-monitor

Supported Platforms	EX Series, M Series, MX Series, PTX Series, QFabric System, QFX Series, T Series
Syntax	show snmp health-monitor <alarms <detail>> <logs>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.
Options	none —Display information about all health monitor alarms and logs. alarms <detail> —(Optional) Display detailed information about health monitor alarms. logs —(Optional) Display information about health monitor logs.
Required Privilege Level	view
List of Sample Output	show snmp health-monitor on page 926 show snmp health-monitor alarms detail on page 928
Output Fields	Table 150 on page 924 describes the output fields for the show snmp health-monitor command. Output fields are listed in the approximate order in which they appear.

Table 150: show snmp health-monitor Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
Variable description	Description of the health monitor object instance being monitored.	All levels
Variable name	Name of the health monitor object instance being monitored.	All levels
Value	Current value of the monitored variable in the most recent sample interval.	All levels

Table 150: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <ul style="list-style-type: none"> Alarms: <ul style="list-style-type: none"> active—Entry is fully configured and activated. falling threshold crossed—Value of the variable has crossed the lower threshold limit. rising threshold crossed—Value of the variable has crossed the upper threshold limit. under creation—Entry is being configured and is not yet activated. startup—Alarm is waiting for the first sample of the monitored variable. object not available—Monitored variable of that type is not available to the health monitor agent. instance not available—Monitored variable's instance is not available to the health monitor agent. object type invalid—Monitored variable is not a numeric value. object processing errored—An error occurred when the monitored variable was processed. unknown—State is not one of the above. 	All levels
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of absolute value or delta value .	detail
Startup alarm	<p>Alarm that might be sent when this entry is first activated, depending on the following criteria:</p> <ul style="list-style-type: none"> Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is either falling alarm or rising or falling alarm. Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is rising alarm. Value of the alarm is between the thresholds. 	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.	detail
Creator	Mechanism by which the entry was configured (Health Monitor).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value as a percentage of the maximum possible value.	detail

Table 150: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Falling threshold	Lower limit threshold value as a percentage of the maximum possible value.	detail
Rising event index	Event triggered when the rising threshold is crossed.	detail
Falling event index	Event triggered when the falling threshold is crossed.	detail

Sample Output

show snmp health-monitor

```
user@host> show snmp health-monitor
```

```
Alarm
Index  Variable description                                Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                      58 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                      0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                        0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0                    35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon                                    0 active
      Chassis daemon                                50 active
      Firewall daemon                                0 active
      Interface daemon                               5 active
      SNMP daemon                                    11 active
      MIB2 daemon                                    42 active
      Sonet APS daemon                               0 active
      VRRP daemon                                    0 active
      Alarm daemon                                    3 active
      PFE daemon                                     0 active
      CRAFT daemon                                   0 active
      Traffic sampling control daemon                 0 active
      Ilmi daemon                                    0 active
      Remote operations daemon                       0 active
      CoS daemon                                     0 active
      Pic Services Logging daemon                     0 active
      Internal Routing Service Daemon                 3 active
      Network Access Service daemon                   0 active
      Forwarding UDP daemon                           0 active
      Routing socket proxy daemon                     0 active
      Disk Monitoring daemon                          1 active
      Inet daemon                                    0 active
      Syslog daemon                                   0 active
      Adaptive Services PIC daemon                    0 active
      ECC parity errors logging Daemon                 0 active
      Layer 2 Tunneling Protocol daemon                0 active
```



```

        PPPoE daemon                      3 active
        Redundancy device daemon          0 active
        PPP daemon                        0 active
        Dynamic Flow Capture Daemon       0 active

32776 Health Monitor: jroute daemon CPU utilization
        Routing protocol daemon           1 active
        Management daemon                 0 active
        Management daemon                 0 active
        Command line interface            4 active
        Periodic Packet Management daemon 0 active
        Link Management daemon            0 active
        Pragmatic General Multicast daemon 0 active
        Bidirectional Forwarding Detection daemon 0 active
        SRC daemon                        0 active
        audit daemon                      0 active
        Event daemon                      0 active

32777 Health Monitor: jcrypto daemon CPU utilization
        IPSec Key Management daemon       0 active

32779 Health Monitor: jkernel daemon Memory utilization
        Init daemon                      47384 active
        Chassis daemon                   20204 active
        Firewall daemon                  1956 active
        Interface daemon                 3340 active
        SNMP daemon                      4540 active
        MIB2 daemon                     3880 active
        Sonet APS daemon                 2632 active
        VRRP daemon                     2672 active
        Alarm daemon                     1856 active
        PFE daemon                       2600 active
        CRAFT daemon                     2000 active
        Traffic sampling control daemon   3164 active
        Ilmi daemon                      2132 active
        Remote operations daemon          2964 active
        CoS daemon                       3044 active
        Pic Services Logging daemon       1944 active
        Internal Routing Service Daemon   1392 active
        Network Access Service daemon     1992 active
        Forwarding UDP daemon             1876 active
        Routing socket proxy daemon        1296 active
        Disk Monitoring daemon            1180 active
        Inet daemon                      1296 active
        Syslog daemon                    1180 active
        Adaptive Services PIC daemon      3220 active
        ECC parity errors logging Daemon  1100 active
        Layer 2 Tunneling Protocol daemon 3372 active
        PPPoE daemon                     1424 active
        Redundancy device daemon          1820 active
        PPP daemon                       2060 active
        Dynamic Flow Capture Daemon       10740 active

32780 Health Monitor: jroute daemon Memory utilization
        Routing protocol daemon           8104 active
        Management daemon                 13360 active
        Management daemon                 19252 active
        Command line interface            9912 active
        Periodic Packet Management daemon 1484 active
        Link Management daemon            2016 active
        Pragmatic General Multicast daemon 1968 active
        Bidirectional Forwarding Detection daemon 1956 active

```

SRC daemon	1772 active
audit daemon	1772 active
Event daemon	1808 active

32781 Health Monitor: jcrypto daemon Memory utilization	
IPSec Key Management daemon	5600 active

show snmp health-monitor alarms detail

```
user@host> show snmp health-monitor alarms detail
```

Alarm Index 32768:

Variable name	jnxHrStoragePercentUsed.1
Variable OID	1.3.6.1.4.1.2636.3.31.1.1.1.1.1
Sample type	absolute value
Startup alarm	rising alarm
Owner	Health Monitor: root file system utilization
Creator	Health Monitor
State	active
Sample interval	300 seconds
Rising threshold	80
Falling threshold	70
Rising event index	32768
Falling event index	32768
Instance Value:	58
Instance State:	active

Alarm Index 32769:

Variable name	jnxHrStoragePercentUsed.2
Variable OID	1.3.6.1.4.1.2636.3.31.1.1.1.1.2
Sample type	absolute value
Startup alarm	rising alarm
Owner	Health Monitor: /config file system utilization
Creator	Health Monitor
State	active
Sample interval	300 seconds
Rising threshold	80
Falling threshold	70
Rising event index	32768
Falling event index	32768
Instance Value:	0
Instance State:	active

Alarm Index 32770:

Variable name	jnxOperatingCPU.9.1.0.0
Variable OID	1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
Sample type	absolute value
Startup alarm	rising alarm
Owner	Health Monitor: RE 0 CPU utilization
Creator	Health Monitor
State	active
Sample interval	300 seconds
Rising threshold	80
Falling threshold	70
Rising event index	32768
Falling event index	32768
Instance Value:	0

Instance State: active

Alarm Index 32773:

Variable name jnxOperatingBuffer.9.1.0.0
 Variable OID 1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0
 Sample type absolute value
 Startup alarm rising alarm
 Owner Health Monitor: RE 0 Memory utilization

Creator Health Monitor
 State active
 Sample interval 300 seconds
 Rising threshold 80
 Falling threshold 70
 Rising event index 32768
 Falling event index 32768
 Instance Value: 35
 Instance State: active

Alarm Index 32775:

Variable name sysAppElmtRunCPU.3
 Variable OID 1.3.6.1.2.1.54.1.2.3.1.9.3
 Sample type delta value
 Startup alarm rising alarm
 Owner Health Monitor: jkernel daemon CPU utilization

Creator Health Monitor
 State active
 Sample interval 300 seconds
 Rising threshold 24000
 Falling threshold 21000
 Rising event index 32768
 Falling event index 32768

Instance Name: sysAppElmtRunCPU.3.1.1
 Instance Description: Init daemon
 Instance Value: 0
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.2.2786
 Instance Description: Chassis daemon
 Instance Value: 50
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.3.2938
 Instance Description: Firewall daemon
 Instance Value: 0
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.4.2942
 Instance Description: Interface daemon
 Instance Value: 5
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.7.7332
 Instance Description: SNMP daemon
 Instance Value: 11
 Instance State: active

Instance Name: sysAppElmtRunCPU.3.9.2914
 Instance Description: MIB2 daemon

Instance Value: 42
Instance State: active

Instance Name: sysAppElemRunCPU.3.12.2916
Instance Description: Sonet APS daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.13.2917
Instance Description: VRRP daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.14.2787
Instance Description: Alarm daemon
Instance Value: 3
Instance State: active

Instance Name: sysAppElemRunCPU.3.15.2940
Instance Description: PFE daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.16.2788
Instance Description: CRAFT daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElemRunCPU.3.17.2918
Instance Description: Traffic sampling control daemon

---(more 23%)---

show snmp inform-statistics

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [OCX1100](#), [PTX Series](#), [QFabric System](#), [QFX Series](#), [T Series](#)

Syntax `show snmp inform-statistics`

Release Information Command introduced in Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display information about Simple Network Management Protocol (SNMP) inform requests.

Options This command has no options.

Required Privilege Level view

List of Sample Output [show snmp inform-statistics on page 932](#)

Output Fields [Table 151 on page 931](#) describes the output fields for the `show snmp inform-statistics` command. Output fields are listed in the approximate order in which they appear.

Table 151: show snmp inform-statistics Output Fields

Field Name	Field Description
Target Name	Name of the device configured to receive and respond to SNMP informs.
Address	IP address of the target device.
Sent	Number of informs sent to the target device and acknowledged by the target device.
Pending	Number of informs held in memory pending a response from the target device.
Discarded	Number of informs discarded after the specified number of retransmissions to the target device were attempted.
Timeouts	Number of informs that did not receive an acknowledgement from the target device within the timeout specified.
Probe Failures	Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address).

Sample Output

show snmp inform-statistics

```
user@host> show snmp inform-statistics
Inform Request Statistics:
Target Name: TA1_v3_md5_none Address: 172.17.20.184
Sent: 176, Pending: 0
Discarded: 0, Timeouts: 0, Probe Failures: 0
Target Name: TA2_v3_sha_none Address: 192.168.110.59
Sent: 0, Pending: 4
Discarded: 84, Timeouts: 0, Probe Failures: 258
Target Name: TA5_v2_none Address: 172.17.20.184
Sent: 0, Pending: 0
Discarded: 2, Timeouts: 10, Probe Failures: 0
```

show snmp mib

Supported Platforms [EX Series](#), [M Series](#), [MX Series](#), [OCX1100](#), [PTX Series](#), [QFX Series](#), [T Series](#)

Syntax `show snmp mib (get | get-next | walk) (ascii | decimal) object-id`

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
ascii and **decimal** options introduced in Junos OS Release 9.6.
ascii and **decimal** options introduced in Junos OS Release 9.6 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.
 Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.

Options **get**—Retrieve and display one or more SNMP object values.

get-next—Retrieve and display the next SNMP object values.

walk—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.

ascii—Display the SNMP object's string indices as an ASCII-key representation.

decimal—Display the SNMP object values in the decimal (default) format. The **decimal** option is the default option for this command. Therefore, issuing the **show snmp mib (get | get-next | walk) decimal object-id** and the **show snmp mib (get | get-next | walk) object-id** commands display the same output.

object-id—The object can be represented by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as **interfaces**). When entering multiple objects, enclose the objects in quotation marks.

Required Privilege Level **snmp**—To view this statement in the configuration.

List of Sample Output [show snmp mib get on page 934](#)
[show snmp mib get \(Multiple Objects\) on page 934](#)
[show snmp mib get \(Layer 2 Policer\) on page 934](#)
[show snmp mib get-next on page 934](#)
[show snmp mib get-next \(Specify an OID\) on page 934](#)
[show snmp mib walk on page 934](#)
[show snmp mib walk \(QFX Series\) on page 935](#)
[show snmp mib walk decimal on page 935](#)
[show snmp mib walk \(ASCII\) on page 935](#)
[show snmp mib walk \(Multiple Indices\) on page 935](#)

[show snmp mib walk decimal \(Multiple Indices\) on page 935](#)

Output Fields [Table 152 on page 934](#) describes the output fields for the **show snmp mib** command. Output fields are listed in the approximate order in which they appear.

Table 152: show snmp mib Output Fields

Field Name	Field Description
<i>name</i>	Object name and numeric instance value.
<i>object value</i>	Object value. The Junos OS translates OIDs into the corresponding object names.

Sample Output

show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get "sysObjectID.0 sysUpTime.0"
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
Networks, Inc.
```

show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004
Juniper Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```


show snmp mib walk (QFX Series)

```

user@switch> show snmp mib walk system
sysDescr.0      = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS
11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0   = jnxProductQFX3500
sysUpTime.0     = 138980301
sysContact.0    = System Contact
sysName.0       = LabQFX3500
sysLocation.0   = Lab
sysServices.0   = 4

```

show snmp mib walk decimal

```

user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100

```

show snmp mib walk (ASCII)

```

show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100

```

show snmp mib walk (Multiple Indices)

```

show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....

```

show snmp mib walk decimal (Multiple Indices)

```

show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....

```

show snmp rmon

Supported Platforms	EX Series, M Series, MX Series, PTX Series, T Series
Syntax	<code>show snmp rmon</code> <code><alarms <brief detail> events <brief detail> logs></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms and events.
Options	none —Display information about all RMON alarms and events. alarms —(Optional) Display information about RMON alarms. brief detail —(Optional) Display brief or detailed information about RMON alarms or events. events —(Optional) Display information about RMON events. logs —(Optional) Display information about RMON monitoring logs.
Required Privilege Level	view
List of Sample Output	show snmp rmon on page 938 show snmp rmon alarms detail on page 938 show snmp rmon events detail on page 939
Output Fields	Table 153 on page 936 describes the output fields for the show snmp rmon command. Output fields are listed in the approximate order in which they appear.

Table 153: show snmp rmon Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels

Table 153: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <p>Alarms:</p> <ul style="list-style-type: none"> • active—Entry is fully configured and activated. • falling threshold crossed—Value of the variable has crossed the lower threshold limit. • rising threshold crossed—Value of the variable has crossed the upper threshold limit. • under creation—Entry is being configured and is not yet activated. • startup—Alarm is waiting for the first sample of the monitored variable. • object not available—Monitored variable of that type is not available to the SNMP agent. • instance not available—Monitored variable's instance is not available to the SNMP agent. • object type invalid—Monitored variable is not a numeric value. • object processing errored—An error occurred when the monitored variable was processed. • unknown—State is not one of the above. <p>Events:</p> <ul style="list-style-type: none"> • active—Entry has been fully configured and activated. • under creation—Entry is being configured and is not yet activated. • unknown—State is not one of the above. 	All levels
Variable name	Name of the SNMP object instance being monitored.	All levels
Event Index	Event identifier.	All levels
Type	<p>Type of notification made when an event is triggered. It can be one of the following:</p> <ul style="list-style-type: none"> • log—A system log message is generated and an entry is made to the log table. • snmptrap—An SNMP trap is sent to the configured destination. • log and trap—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination. • none—Neither log nor trap will be sent. 	detail
Last Event	Date and time of the last event. It has the format <i>yyyy-mm-dd hh:mm:ss timezone</i> .	brief
Community	Identifies the trap group used for sending the SNMP trap.	detail
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of absolute value or delta value .	detail

Table 153: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria: <ul style="list-style-type: none"> Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is either falling alarm or rising or falling alarm. Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is rising alarm. Value of the alarm is between the thresholds. 	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.	detail
Creator	Mechanism by which the entry was configured (CLI or SNMP).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value configured by the user.	detail
Falling threshold	Lower limit threshold value configured by the user.	detail
Rising event index	Event triggered when the rising threshold is crossed.	detail
Falling event index	Event triggered when the falling threshold is crossed.	detail
Current value	Current value of the monitored variable in the most recent sample interval.	detail

Sample Output

show snmp rmon

```

user@host> show snmp rmon
Alarm
Index  State                      Variable name
      1  falling threshold crossed  ifInOctets.1

Event
Index  Type                      Last Event
      1  log and trap              2002-01-30 01:13:01 PST

```

show snmp rmon alarms detail

```

user@host> show snmp rmon alarms detail

```

```
Alarm Index 1:
Variable name      ifInOctets.1
Variable OID       1.3.6.1.2.1.2.2.1.10.1
Sample type        delta value
Startup alarm      rising or falling alarm
Owner              monitor
Creator            CLI
State              falling threshold crossed
Sample interval    60 seconds
Rising threshold   100000
Falling threshold  80000
Rising event index 1
Falling event index 1
Current value      0
```

show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
Type              log and trap
Community          boy-elroy
Last event         2002-01-30 01:13:01 PST
Creator            CLI
State              active
```

show snmp statistics

Supported Platforms	EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, T Series
Syntax	show snmp statistics <subagents>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for OCX Series switches. Option subagents introduced in Junos OS Release 14.2.
Description	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
Options	subagents —(Optional) Display the statistics of the protocol data unit (PDU), the number of SNMP requests and responses per subagent, and the SNMP statistics received from each subagent per logical system.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear snmp statistics on page 835
List of Sample Output	show snmp statistics on page 945 show snmp statistics subagents on page 945
Output Fields	Table 154 on page 941 describes the output fields for the show snmp statistics command. Output fields are listed in the approximate order in which they appear.

Table 154: show snmp statistics Output Fields

Field Name	Field Description
Input	<p>Information about received packets:</p> <ul style="list-style-type: none"> • Packets(snmplnPkts)—Total number of messages delivered to the SNMP entity from the transport service. • Bad versions—(snmplnBadVersions) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • Bad community names—(snmplnBadCommunityNames) Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • Bad community uses—(snmplnBadCommunityUses) Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • ASN parse errors—(snmplnASNParseErrs) Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • Too big—(snmplnTooBigs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of tooBig. • No such names—(snmplnNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmplnBadValues) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of badValue. • Read onlys—(snmplnReadOnlys) Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of readOnly. Only incorrect implementations of SNMP generate this error.

Table 154: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> • General errors—(snmpInGenErrs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of genErr. • Total requests varbinds—(snmpInTotalReqVars) Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP GetRequest and GetNext PDUs. • Total set varbinds—(snmpInSetVars) Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP SetRequest PDUs. • Get requests—(snmpInGetRequests) Total number of SNMP GetRequest PDUs that have been accepted and processed by the SNMP entity. • Get nexts—(snmpInGetNexts) Total number of SNMP GetNext PDUs that have been accepted and processed by the SNMP entity. • Set requests—(snmpInSetRequests) Total number of SNMP SetRequest PDUs that have been accepted and processed by the SNMP entity. • Get responses—(snmpInGetResponses) Total number of SNMP GetResponse PDUs that have been accepted and processed by the SNMP entity. • Traps—(snmpInTraps) Total number of SNMP traps generated by the SNMP entity. • Silent drops—(snmpSilentDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • Proxy drops—(snmpProxyDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned. • Commit pending drops—Number of SNMP packets for Set requests dropped because of a previous pending SNMP Set request on the committed configuration. • Throttle drops—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 154: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> • Unknown security models—(snmpUnknownSecurityModels) Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • Invalid messages—(snmpInvalidMsgs) Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • Unknown pdu handlers—(snmpUnknownPDUHandlers) Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • Unavailable contexts—(snmpUnavailableContexts) Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • Unknown contexts—(snmpUnknownContexts) Total number of requests received for a context that is unknown to the SNMP engine. • Unsupported security levels—(usmStatsUnsupportedSecLevels) Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • Not in time windows—(usmStatsNotInTimeWindows) Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window. • Unknown user names—(usmStatsUnknownUserNames) Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • Unknown engine ids—(usmStatsUnknownEngineIDs) Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine. • Wrong digests—(usmStatsWrongDigests) Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value. • Decryption errors—(usmStatsDecryptionErrors) Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.

Table 154: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Output	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> • Packets—(snmpOutPkts) Total number of messages passed from the SNMP entity to the transport service. • Too big—(snmpOutTooBigs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of tooBig. • No such names—(snmpOutNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmpOutBadValues) Total number of SNMP PDUs generated by the SNMP entity with an error status field of badValue. • General errors—(snmpOutGenErrs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of genErr. • Get requests—(snmpOutGetRequests) Total number of SNMP GetRequest PDUs generated by the SNMP entity. • Get nexts—(snmpOutGetNexts) Total number of SNMP GetNext PDUs generated by the SNMP entity. • Set requests—(snmpOutSetRequests) Total number of SNMP SetRequest PDUs generated by the SNMP entity. • Get responses—(snmpOutGetResponses) Total number of SNMP GetResponse PDUs generated by the SNMP entity. • Traps—(snmpOutTraps) Total number of SNMP traps generated by the SNMP entity.

Table 155 on page 944 describes the output fields for the **show snmp statistics subagents** command. Output fields are listed in the approximate order in which they appear.

Table 155: show snmp statistics subagents Output Fields

Field Name	Field Description
Subagent	Location of the SNMP subagent.
Request PDUs	Number of PDUs requested by the SNMP manager.
Response PDUs	Number of response PDUs sent by the SNMP subagent.
Request Variables	Number of variable bindings on the PDUs requested by the SNMP manager.
Response Variables	Number of variable bindings on the PDUs sent by the SNMP subagent.
Average Response Time	Average time taken by the SNMP subagent to send statistics response.
Maximum Response Time	Maximum time taken by the SNMP subagent to send the statistics response.

Sample Output

show snmp statistics

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too bigs: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0
```

show snmp statistics subagents

```
user@host> show snmp statistics subagents

Subagent: /var/run/cosd-20
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/pfed-30
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/rmopd-15
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
  Maximum Response Time(ms): 0.00

Subagent: /var/run/chassisd-30
  Request PDUs: 33116, Response PDUs: 33116,
  Request Variables: 33116, Response Variables: 33116,
  Average Response Time(ms): 1.83,
  Maximum Response Time(ms): 203.48

Subagent: /var/run/pkid-13
  Request PDUs: 0, Response PDUs: 0,
  Request Variables: 0, Response Variables: 0,
  Average Response Time(ms): 0.00,
```

Maximum Response Time(ms): 0.00

Subagent: /var/run/apsd-13
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/dfcd-32
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/mib2d-33
Request PDUs: 74211, Response PDUs: 74211,
Request Variables: 74211, Response Variables: 74211,
Average Response Time(ms): 2.30,
Maximum Response Time(ms): 51.04

Subagent: /var/run/license-check-16
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/craftd-14
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/bfdd-19
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/smihelperd-24
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/cfmd-18
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/rpd_snmp
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

Subagent: /var/run/l2tpd-18
Request PDUs: 0, Response PDUs: 0,
Request Variables: 0, Response Variables: 0,
Average Response Time(ms): 0.00,
Maximum Response Time(ms): 0.00

show snmp stats-response-statistics

Supported Platforms	ACX Series, E Series, EX Series, M Series, MX Series, PTX Series, QFX Series, T Series
Syntax	show snmp stats-response-statistics
Release Information	Command introduced in Junos OS Release 14.2.
Description	Display statistics of SNMP statistics responses sent from the Packet Forwarding Engine during the MIB II process (mib2d).
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show snmp stats-response-statistics on page 949
Output Fields	Table 156 on page 948 describes the output fields for the show snmp stats-response-statistics command. Output fields are listed in the approximate order in which they appear.

Table 156: show snmp stats-response-statistics Output Fields

Field Name	Field Description
Average response time statistics	<p>Display the average response time in milliseconds per protocol data unit (PDU) by snmpd. It includes the following information:</p> <ul style="list-style-type: none"> • Stats Type—Type of SNMP statistics. • Stats Responses—Number of SNMP statistics responses received from the Packet Forwarding Engine. • Average Response Time—Average time taken to receive the statistics response from the Packet Forwarding Engine in milliseconds.
Bucket statistics	<p>Information about SNMP statistics responses:</p> <ul style="list-style-type: none"> • Bucket Type—Category of time intervals in which SNMP statistics responses are received from the Packet Forwarding Engine. • Stats Responses—Number of SNMP statistics responses received from the Packet Forwarding Engine.
Bad responses	<p>Information about top 20 bad responses from a subagent:</p> <ul style="list-style-type: none"> • Response—Time taken to receive the SNMP statistics response from the Packet Forwarding Engine in milliseconds. • Request Time—Date and time of SNMP request. • Key—Display the attribute of SNMP Stats Type. For example, in the case of SNMP statistics responses for interfaces, the Key value is SNMP ifIndex, and for firewalls, the Key value is the filter name.

Sample Output

show snmp stats-response-statistics

```
user@host> show snmp stats-response-statistics
```

Average response time statistics:

Stats Type	Stats Responses	Average Response Time (ms)
ifd(non ae)	34182	175.48
ifd(ae)	0	0.00
ifl(non ae)	5472	5.40
ifl(ae)	0	0.00
firewall	15	1141.73

Bucket statistics:

Bucket Type(ms)	Stats Responses
0 - 10	39078
11 - 50	588
51 - 100	0
101 - 200	0
201 - 500	1
501 - 1000	2
1001 - 2000	0
2001 - 5000	0
More than 5001	0

Bad responses:

Response (ms)	Request Time (UTC)	Stats Type	Key
927.80	2014-03-26 05:44:16	firewall	__default_arp_policer__
908.68	2014-03-26 05:44:16	firewall	__default_bpdu_filter__
421.00	2014-03-26 05:46:25	ifd(non ae)	504
49.76	2014-04-13 04:15:18	ifd(non ae)	503
49.62	2014-04-13 04:30:18	ifd(non ae)	504
48.52	2014-04-05 10:06:55	ifd(non ae)	504
47.61	2014-04-11 04:06:27	ifd(non ae)	505
47.38	2014-04-13 03:30:18	ifd(non ae)	501
47.22	2014-03-27 20:08:07	ifd(non ae)	502
46.26	2014-03-31 13:08:58	ifd(non ae)	506
46.00	2014-04-13 04:00:18	ifd(non ae)	503
45.95	2014-04-05 17:15:17	ifd(non ae)	503
45.75	2014-04-15 13:06:10	ifd(non ae)	507
45.60	2014-04-01 03:07:28	ifd(non ae)	517
45.56	2014-04-08 13:09:15	ifd(non ae)	502
45.23	2014-04-13 03:15:18	ifd(non ae)	501
45.15	2014-04-05 16:45:17	ifd(non ae)	501
44.74	2014-04-08 22:08:47	ifd(non ae)	505
44.10	2014-04-05 16:30:17	ifd(non ae)	501
44.00	2014-04-08 09:09:23	ifd(non ae)	524

show snmp v3

Supported Platforms	EX Series, M Series, MX Series, PTX Series, T Series
Syntax	<code>show snmp v3</code> <code><access <brief detail> community general groups notify <filter> target <address parameters> users></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.
Options	none —Display all of the SNMPv3 operating configuration. access —(Optional) Display SNMPv3 access information. brief detail —(Optional) Display brief or detailed information about SNMPv3 access information. community —(Optional) Display SNMPv3 community information. general —(Optional) Display SNMPv3 general information. groups —(Optional) Display SNMPv3 security-to-group information. notify <filter> —(Optional) Display SNMPv3 notify and, optionally, notify filter information. target <address parameters> —(Optional) Display SNMPv3 target and, optionally, either target address or target parameter information. users —(Optional) Display SNMPv3 user information.
Additional Information	To edit the default display of the show snmp v3 command, specify options in the show statement at the [edit snmp v3] hierarchy level.
Required Privilege Level	view
List of Sample Output	show snmp v3 on page 952
Output Fields	Table 157 on page 951 describes the output fields for the show snmp v3 command. Output fields are listed in the approximate order in which they appear.

Table 157: show snmp v3 Output Fields

Field Name	Field Description
Access control	<p>Information about access control:</p> <ul style="list-style-type: none"> • Group—Group name for which the configured access privileges apply. The group, together with the context prefix and the security model and security level, forms the index for this table. • Context prefix—SNMPv3 context for which the configured access privileges apply. • Security model/level—Security model and security level for which the configuration access privileges apply. • Read view—Identifies the MIB view applied to SNMPv3 read operations. • Write view—Identifies the MIB view applied to SNMPv3 write operations. • Notify view—Identifies the MIB view applied to outbound SNMP notifications.
Engine	<p>Information about local engine configuration:</p> <ul style="list-style-type: none"> • Local engine ID—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • Engine boots—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • Engine time—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • Max msg size—Maximum message size the sender can accommodate.
Engine ID	<p>Information about engine ID:</p> <ul style="list-style-type: none"> • Local engine ID—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • Engine boots—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • Engine time—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • Max msg size—Maximum message size the sender can accommodate. • Engine ID—SNMPv3 engine ID associated with each user. • User—SNMPv3 user. • Auth/Priv—Authentication and encryption algorithm available for use by each user. • Storage—Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status. • Status—Status of the conceptual row. Only rows with an active status are used by the SNMPv3 engine.
Group name	Name of the group to which this entry belongs.
Security model	Identifies the security model context for the security name.
Security name	Used with the security model; identifies a specific security name instance. Each security model/security name combination can be assigned to a specific group.
Storage type	Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.
Status	Status of the conceptual row. Only rows with active status are used by the SNMPv3 engine.

Sample Output

show snmp v3

```
user@host> show snmp v3
Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots:      38
Engine time:       64583 seconds
Max msg size:      2048 bytes

Engine ID: local
  User          Auth/Priv  Storage  Status
  user1         md5/des   nonvolatile active
  user2         sha/none  nonvolatile active
  user3         none/none nonvolatile active

Engine ID: 81 00 0a 4c 04 64 64 64 64
  User          Auth/Priv  Storage  Status
  UNEW         md5/none  nonvolatile active
Group name      Security  Security  Storage  Status
               model    name      type
g1             usm      user1     nonvolatile active
g2             usm      user2     nonvolatile active
g3             usm      user3     nonvolatile active

Access control:
Group          Context Security  Read  Write  Notify
               prefix model/level view  view  view
g1             usm/privacy v1    v1
g2             usm/authent v1    v1
g3             usm/none   v1    v1
```

show system alarms

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

Syntax `show system alarms`

Release Information Command introduced in Junos OS Release 11.1 for SRX Series devices.

Description Display active system alarms.

Options This command has no options.

Additional Information System alarms are preset. You cannot modify them, although you can configure them to appear automatically in the J-Web user interface or CLI. They include a **configuration** alarm that appears when no rescue configuration alarm is set and a **license** alarm that appears when a software feature is configured but no valid license is configured for the feature.

Required Privilege Level admin

List of Sample Output [show system alarms on page 953](#)

Sample Output

`show system alarms`

```
user@host> show system alarms
5 alarms currently active
Alarm time      Class  Description
2012-05-29 16:47:18 UTC Major  /var partition usage crossed critical threshold
2012-05-29 16:47:18 UTC Minor  /var partition usage crossed high threshold
2012-05-29 16:47:18 UTC Major  /root partition usage crossed critical threshold
2012-05-29 16:47:18 UTC Minor  /root partition usage crossed high threshold
2012-05-29 16:47:18 UTC Minor  Rescue configuration is not set
```

show system alarms

Supported Platforms EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series, T Series

Syntax show system alarms

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display active system alarms.

Options This command has no options.

Additional Information System alarms are preset. You cannot modify them, although you can configure them to appear automatically in the J-Web user interface or CLI. They include a *configuration* alarm that appears when no rescue configuration alarm is set and a *license* alarm that appears when a software feature is configured and no valid license is configured for the feature. You can also determine when a license will expire from syslog messages that appear starting from four weeks before expiry of the license. On EX6200 switches, an alarm can be triggered by an internal link error.

The the logic for multiple feature licenses is based on the highest validity among the licenses. Also, for capacity non-cumulative, exclusive type licenses (such as for scale), the logic is based on the highest validity of the license.



NOTE: As of Junos OS Release 17.3R1, the logic for multiple capacity type licenses and when their expiry raises alarms was changed. Before, the behavior had alarms and syslog messages for expiring licenses raised based on the highest validity, which would mislead users in the case of a license expiring earlier than the highest validity license. The newer behavior has the about-to-expire logic based on the first expiring license.

For more information about system alarms, see the *Junos OS Administration Library*.

In Junos OS release 11.1 and later, alarms for fans also show the slot number of the malfunctioning fans in the CLI output.

Starting with Junos OS Release 13.2, you can view degraded fabric alarms on a routing matrix based on TX Matrix Plus router with 3D SIBs. The alarm indicates that the source FPC is running with a degraded fabric condition. This alarm is an early warning of a possible fabric black-hole condition. When the degraded fabric alarm is raised on the source FPC, you can take remedial action to avoid a fabric black-hole condition. The degraded fabric alarm is raised on the source FPC if both the following conditions are met:

- The active Packet Forwarding Engine destinations are reachable on one or no active switching planes.
- At least one of the inactive switching planes has a fault that causes the destination Packet Forwarding Engine to become unreachable.

Required Privilege Level view

Related Documentation

- *show chassis alarms*

List of Sample Output [show system alarms on page 955](#)
[show system alarms \(Fan Tray\) on page 955](#)
[show system alarms \(QFX Series and OCX Series\) on page 956](#)
[show system alarms \(EX6200\) on page 956](#)
[show system alarms \(TX Matrix Plus router with 3D SIBs\) on page 956](#)

Output Fields [Table 158 on page 955](#) lists the output fields for the **show system alarms** command. Output fields are listed in the approximate order in which they appear.

Table 158: show system alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: Minor or Major .
Description	Information about the alarm.

Sample Output

show system alarms

```
user@host> show system alarms
2 alarms currently active
Alarm time          Class    Description
2005-02-24 17:29:34 UTC  Minor    IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC  Minor    Rescue configuration is not sent
```

show system alarms (Fan Tray)

```
user@host> show system alarms
4 alarms currently active
Alarm time          Class    Description
2010-11-11 20:27:38 UTC  Major    Side Fan Tray 7 Failure
2010-11-11 20:27:13 UTC  Minor    Side Fan Tray 7 Overspeed
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 5 Failure
2010-11-11 20:27:13 UTC  Major    Side Fan Tray 0 Failure
```

show system alarms (QFX Series and OCX Series)

```
user@switch> show system alarms
2 alarms currently active
Alarm time Class Description
2005-02-24 17:29:34 UTC Minor Rescue configuration is not sent
```

show system alarms (EX6200)

```
user@switch> show system alarms
2 alarms currently active
Alarm time          Class Description
2013-04-05 16:51:41 PDT Major FPC 8 internal link errors detected
2013-04-04 18:05:35 PDT Minor Rescue configuration is not set
```

show system alarms (TX Matrix Plus router with 3D SIBs)

```
user@router> show system alarms

sfc0-re0:
-----
2 alarms currently active
Alarm time          Class Description
2013-05-08 18:13:58 UTC Major LCC 0 Major Errors
2013-05-08 17:48:46 UTC Major LCC 7 Major Errors

lcc0-re1:
-----
1 alarm currently active
Alarm time          Class Description
2013-05-08 18:19:24 UTC Major FPC 1 degraded fabric condition detected

lcc7-re0:
-----
1 alarm currently active
Alarm time          Class Description
2013-05-08 18:19:24 UTC Major FPC 7 degraded fabric condition detected
```

show system resource-monitor fpc

Supported Platforms EX9200, EX9204, EX9208, EX9214, MX104, MX2010, MX2020, MX240, MX480, MX80, MX960

Syntax show system resource-monitor fpc *slot-number*

Release Information Command introduced in Junos OS Release 15.1 for MX80, MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.
Command introduced in Junos OS Release 16.1 for EX9200 switches.

Description Display the utilization of memory resources on the Packet Forwarding Engines of an FPC. The filter memory denotes the filter counter memory used for firewall filter counters. The asterisk (*) displayed next to each of the memory regions denotes the ones for which the configured threshold is being currently exceeded.

Options **fpc slot-number**—Display the Junos OS utilization information of memory resources for the specified slot number in which the FPC is installed.

- MX80 router—Replace **fpc-slot** with a value from 1. This command is not supported on FPC slot 0.
- MX104 router—Replace **fpc-slot** with a value from 0 through 2.
- MX240 router—Replace **fpc-slot** with a value from 0 through 2.
- MX480 router—Replace **fpc-slot** with a value from 0 through 5.
- MX-960 router—Replace **fpc-slot** with a value from 0 through 11.
- MX2010 router—Replace **fpc-slot-number** with a value from 0 through 9.
- MX2020 router—Replace **fpc-slot-number** with a value from 0 through 19.
- EX9204 switches—Replace **fpc-slot-number** with a value from 0 through 2.
- EX9208 switches—Replace **fpc-slot-number** with a value from 0 through 5.
- EX9214 switches—Replace **fpc-slot-number** with a value from 0 through 11.

Additional Information The filter memory denotes the filter counter memory used for firewall filter counters. From the Ukern perspective, MPC5E contains only one Packet Forwarding Engine instance. The **show chassis fabric plane** command output displays the state of fabric plane connections to the Packet Forwarding Engine. Because two Packet Forwarding Engines exist, you notice PFE-0 and PFE-1 in the output.

Required Privilege Level view

- Related Documentation**
- *show system resource-monitor ifd-cos-queue-mapping fpc*
 - *show system resource-monitor subscribers-limit*

List of Sample Output [show system resource-monitor fpc on page 958](#)

Output Fields [Table 159 on page 958](#) lists the output fields for the **show system resource-monitor fpc** command. Output fields are listed in the approximate order in which they appear.

Table 159: show system resource-monitor fpc Output Fields

Field Name	Field Description
Free Heap Memory Watermark	Configured watermark value for the percentage of free memory space used for ukernel or heap memory to be monitored
Free FW Memory Watermark	Configured watermark value for the percentage of free memory space used for firewall or filter memory to be monitored
Free NH Memory Watermark	Configured watermark value for the percentage of free memory space used for next-hop memory to be monitored
* - watermark reached	An asterisk (*) displayed beside any of the memory regions denotes the memory types for which the configured threshold is being currently exceeded.
Slot #	Slot number in which the line card is installed
PFE #	Number or identifier of the Packet Forwarding Engine in the specified line card slot
Heap % free	Percentage of free space associated with heap or ukernel memory
Encap mem % free	Percentage of free space associated with encapsulation memory
NH mem % free	Percentage of free space associated with next-hop memory
Filter / FW mem % free	Percentage of free space associated with firewall or filter memory

Sample Output

show system resource-monitor fpc

```
user@host> show system resource-monitor fpc
FPC Resource Usage Summary
```



```

Throttle                : Enabled
Heap Mem Threshold      : 80 %
Round Trip Delay Threshold :120 ms
IFL Counter Threshold   : 80 %
Filter Counter Threshold : 80 %
Expansion Threshold     : 95 %
MFS threshold           : 80 %           Used : 1
Slot # 1
  Client allowed        : Yes
  Service allowed       : Yes
  Client denied         : 0             Performance Denial Client : 0
  Service denied        : 0             Performance Denial Service :0
  Heap memory used      : 183985808     In % : 9
  Average Round-trip Delay : 127 ms
  Round Trip Delay      : 130 ms *      MAX session rate allowed(%) : 80
    Filter memory      IFL memory      Expansion memory
    PFE #      used | %      used | %      used | %
    0      29696      0      5056      0      0      0
    1      29536      0      4896      0      0      0
Slot # 2
  Client allowed        : Yes
  Service allowed       : Yes
  Client denied         : 0             Performance Denial Client : 0
  Service denied        : 0             Performance Denial Service :0
  Heap memory used      : 183982960     In % : 9
  Average Round-trip Delay : 98 ms
  Round Trip Delay      : 100 ms        MAX session rate allowed(%) : 100
    Filter memory      IFL memory      Expansion memory
    PFE #      used | %      used | %      used | %
    0      29856      0      5216      0      0      0
    1      29376      0      4736      0      0      0

```

