



Junos[®] OS

Network Monitoring and Troubleshooting Guide for Security Devices

Release

15.1X49-D40



Modified: 2016-10-04

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Junos[®] OS Network Monitoring and Troubleshooting Guide for Security Devices
15.1X49-D40
Copyright © 2016, 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.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xix
	Documentation and Release Notes	xix
	Supported Platforms	xix
	Using the Examples in This Manual	xix
	Merging a Full Example	xx
	Merging a Snippet	xx
	Documentation Conventions	xxi
	Documentation Feedback	xxiii
	Requesting Technical Support	xxiii
	Self-Help Online Tools and Resources	xxiii
	Opening a Case with JTAC	xxiv
Part 1	Overview	
Chapter 1	Introduction to Network Monitoring	3
	Monitoring Overview	3
	Diagnostic Tools Overview	4
	J-Web Diagnostic Tools	4
	CLI Diagnostic Commands	5
Chapter 2	Accounting Options, Source Class Usage, and Destination Class Usage Overview	7
	Accounting Options Overview	7
	Understanding Source Class Usage and Destination Class Usage Options	8
Chapter 3	Gathering Statistics for Accounting Purposes	11
	Accounting Options Configuration	11
	Accounting Options—Full Configuration	11
	Minimum Accounting Options Configuration	13
	Configuring Accounting-Data Log Files	14
	Configuring the Storage Location of the File	15
	Configuring the Maximum Size of the File	16
	Configuring the Maximum Number of Files	16
	Configuring the Start Time for File Transfer	16
	Configuring the Transfer Interval of the File	16
	Configuring Archive Sites	17
	Configuring the Interface Profile	17
	Configuring Fields	18
	Configuring the File Information	18
	Configuring the Interval	19
	Example: Configuring the Interface Profile	19

Configuring the Filter Profile	20
Configuring the Counters	21
Configuring the File Information	21
Configuring the Interval	21
Example: Configuring a Filter Profile	22
Example: Configuring Interface-Specific Firewall Counters and Filter Profiles . . .	23
Configuring SCU or DCU	24
Creating Prefix Route Filters in a Policy Statement	24
Applying the Policy to the Forwarding Table	25
Enabling Accounting on Inbound and Outbound Interfaces	25
Configuring SCU on a Virtual Loopback Tunnel Interface	26
Example: Configuring a Virtual Loopback Tunnel Interface on a Provider Edge Router Equipped with a Tunnel PIC	26
Example: Mapping the VRF Instance Type to the Virtual Loopback Tunnel Interface	27
Example: Sending Traffic Received from the Virtual Loopback Interface Out the Source Class Output Interface	27
Configuring Class Usage Profiles	28
Configuring a Class Usage Profile	28
Configuring the File Information	28
Configuring the Interval	29
Creating a Class Usage Profile to Collect Source Class Usage Statistics . . .	29
Creating a Class Usage Profile to Collect Destination Class Usage Statistics	29
Configuring the MIB Profile	30
Configuring the File Information	30
Configuring the Interval	31
Configuring the MIB Operation	31
Configuring MIB Object Names	31
Example: Configuring a MIB Profile	31
Configuring the Routing Engine Profile	32
Configuring Fields	32
Configuring the File Information	33
Configuring the Interval	33
Example: Configuring a Routing Engine Profile	33

Part 2

Chapter 4

Configuring Monitoring Options

Configuring Interface Alarms	37
Alarm Overview	37
Alarm Types	37
Alarm Severity	38
Alarm Conditions	38
Interface Alarm Conditions	39
System Alarm Conditions	42
Example: Configuring Interface Alarms	43
Monitoring Active Alarms on a Device	46
Monitoring Alarms	47

Chapter 5	Using RPM to Measure Network Performance	49
	RPM Overview	49
	RPM Probes	50
	RPM Tests	50
	Probe and Test Intervals	50
	Jitter Measurement with Hardware Timestamping	51
	RPM Statistics	51
	RPM Thresholds and Traps	53
	RPM for BGP Monitoring	53
	IPv6 RPM Probes	53
	Guidelines for Configuring RPM Probes for IPv6	54
	RPM Support for VPN Routing and Forwarding	55
	Example: Configuring Basic RPM Probes	55
	Example: Configuring RPM Using TCP and UDP Probes	59
	Example: Configuring RPM Probes for BGP Monitoring	62
	Directing RPM Probes to Select BGP Devices	65
	Configuring RPM Timestamping	65
	Configuring IPv6 RPM Probes	66
	Tuning RPM Probes	67
	RPM Configuration Options	68
	Monitoring RPM Probes	72
Chapter 6	Configuring IP Monitoring	77
	IP Monitoring Overview	77
	Understanding IP Monitoring Test Parameters	78
	Example: Configuring IP Monitoring on Branch SRX Series Devices	79
	Understanding IP Monitoring Through Redundant Ethernet Interface Link Aggregation Groups	81
	Example: Configuring IP Monitoring on High-End SRX Series Devices	82
	Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring	87
Part 3	Monitoring Common Security Features	
Chapter 7	Displaying Real-Time Information from Device to Host	95
	Displaying Multicast Path Information	95
	Displaying Real-Time Monitoring Information	97
Chapter 8	Monitoring Application Layer Gateways Features	101
	Monitoring H.323 ALG Information	101
	Monitoring MGCP ALGs	102
	Monitoring MGCP ALG Calls	103
	Monitoring MGCP ALG Counters	103
	Monitoring MGCP ALG Endpoints	105
	Monitoring SCCP ALGs	105
	Monitoring SCCP ALG Calls	106
	Monitoring SCCP ALG Counters	106
	Monitoring SIP ALGs	108
	Monitoring SIP ALG Calls	108
	Monitoring SIP ALG Counters	109

	Monitoring SIP ALG Rate Information	111
	Monitoring SIP ALG Transactions	112
	Monitoring Voice ALG H.323	112
	Monitoring Voice ALG MGCP	114
	Monitoring Voice ALG SCCP	117
	Monitoring Voice ALG SIP	120
	Monitoring Voice ALG Summary	125
Chapter 9	Monitoring Class of Service	127
	Monitoring Class-of-Service Performance	127
	Monitoring CoS Interfaces	127
	Monitoring CoS Classifiers	128
	Monitoring CoS Value Aliases	129
	Monitoring CoS RED Drop Profiles	130
	Monitoring CoS Forwarding Classes	131
	Monitoring CoS Rewrite Rules	132
	Monitoring CoS Scheduler Maps	133
	Monitoring CoS Classifiers	135
Chapter 10	Monitoring Interfaces and Switching Functions	137
	Displaying Real-Time Interface Information	137
	Monitoring Address Pools	139
	Monitoring Ethernet Switching	140
	Monitoring GVRP	141
	Monitoring Interfaces	142
	Monitoring MPLS Traffic Engineering Information	143
	Monitoring MPLS Interfaces	144
	Monitoring MPLS LSP Information	144
	Monitoring MPLS LSP Statistics	145
	Monitoring RSVP Session Information	146
	Monitoring MPLS RSVP Interfaces Information	148
	Monitoring PPP	149
	Monitoring PPPoE	149
	Monitoring Spanning Tree	153
	Monitoring the WAN Acceleration Interface	154
Chapter 11	Monitoring NAT	155
	Monitoring NAT	155
	Monitoring Source NAT Information	155
	Monitoring Destination NAT Information	161
	Monitoring Static NAT Information	163
	Monitoring Incoming Table Information	164
	Monitoring Interface NAT Port Information	165
Chapter 12	Monitoring Security Policies	167
	Monitoring Policy Statistics	167
	Monitoring Routing Information	168
	Monitoring Route Information	168
	Monitoring RIP Routing Information	170
	Monitoring OSPF Routing Information	171

	Monitoring BGP Routing Information	173
	Monitoring Security Events by Policy	175
	Monitoring Security Features	177
	Monitoring Policies	177
	Checking Policies	180
	Monitoring Screen Counters	183
	Monitoring IDP Status	185
	Monitoring Flow Gate Information	186
	Monitoring Firewall Authentication Table	187
	Monitoring Firewall Authentication History	189
	Monitoring 802.1x	191
Chapter 13	Monitoring Events, Services and System	193
	Monitoring DHCP Client Bindings	193
	Monitoring Events	193
	Monitoring the System	196
	Monitoring System Properties for SRX Series Devices	196
	Monitoring Chassis Information	198
	System Health Management for Branch SRX Series Devices	200
Chapter 14	Monitoring Unified Threat Management Features	203
	Monitoring Antivirus Scan Engine Status	203
	Monitoring Antivirus Scan Results	204
	Monitoring Antivirus Session Status	206
	Monitoring Content Filtering Configurations	207
	Monitoring Reports	207
	Threats Monitoring Report	208
	Traffic Monitoring Report	212
	Monitoring Web Filtering Configurations	214
Chapter 15	Monitoring VPNs	217
	Monitoring VPNs	217
	Monitoring IKE Gateway Information	217
	Monitoring IPsec VPN—Phase I	221
	Monitoring IPsec VPN—Phase II	222
	Monitoring IPsec VPN Information	223
Part 4	Troubleshooting	
Chapter 16	Configuring Data Path Debugging and Trace Options	231
	Understanding Data Path Debugging for SRX Series Devices	231
	Debugging the Data Path (CLI Procedure)	232
	Example: Configuring End-to-End Debugging on a High-End SRX Series Device	233
	Understanding Security Debugging Using Trace Options	237
	Setting Security Trace Options (CLI Procedure)	237
	Displaying Log and Trace Files	239
	Displaying Output for Security Trace Options	239
	Displaying Multicast Trace Operations	240

	Using the J-Web Traceroute Tool	241
	J-Web Traceroute Results and Output Summary	243
	Understanding Flow Debugging Using Trace Options	243
	Setting Flow Debugging Trace Options (CLI Procedure)	244
	Displaying a List of Devices	245
Chapter 17	Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits	247
	MPLS Connection Checking Overview	247
	Configuring Ping MPLS	249
	Using the ping Command	250
	Using the J-Web Ping Host Tool	252
	J-Web Ping Host Results and Output Summary	254
	Using the J-Web Ping MPLS Tool	255
	J-Web Ping MPLS Results and Output Summary	258
	Pinging Layer 2 Circuits	259
	Pinging Layer 2 VPNs	260
	Pinging Layer 3 VPNs	261
	Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs	262
Chapter 18	Using Packet Capture to Analyze Network Traffic	265
	Packet Capture Overview	265
	Packet Capture on Device Interfaces	266
	Firewall Filters for Packet Capture	267
	Packet Capture Files	267
	Analysis of Packet Capture Files	267
	Example: Enabling Packet Capture on a Device	268
	Example: Configuring Packet Capture on an Interface	271
	Example: Configuring a Firewall Filter for Packet Capture	273
	Example: Configuring Packet Capture for Datapath Debugging	275
	Disabling Packet Capture	278
	Deleting Packet Capture Files	279
	Changing Encapsulation on Interfaces with Packet Capture Configured	280
	Displaying Packet Headers	281
	Using the J-Web Packet Capture Tool	285
	J-Web Packet Capture Results and Output Summary	288
Chapter 19	Troubleshooting Security Devices	291
	Recovering the Root Password for SRX Series Devices	291
	Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only)	292
	Troubleshooting the Link Services Interface	293
	Determine Which CoS Components Are Applied to the Constituent Links	293
	Determine What Causes Jitter and Latency on the Multilink Bundle	295
	Determine If LFI and Load Balancing Are Working Correctly	295
	Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device	302
	Troubleshooting Security Policies	302
	Checking a Security Policy Commit Failure	302
	Verifying a Security Policy Commit	303

	Debugging Policy Lookup	303
	Understanding Log Error Messages for Troubleshooting ISSU-Related Problems	304
	Chassisd Process Errors	304
	Kernel State Synchronization	304
	Installation Related Errors	304
	ISSU Support Related Errors	305
	Redundancy Group Failover Errors	305
	Initial Validation Checks Fail	305
Part 5	Configuration Statements and Operational Commands	
Chapter 20	Configuration Statements	309
	Accounting-Options Configuration Statement Hierarchy	310
	[edit security alarms] Hierarchy Level	312
	[edit security datapath-debug] Hierarchy Level	313
	[edit security traceoptions] Hierarchy Level	314
	accounting-options	314
	action-profile	315
	archive-sites	316
	capture-file (Security)	317
	class-usage-profile	318
	cluster (Chassis)	319
	counters	320
	datapath-debug	321
	decryption-failures	322
	destination-classes	323
	destination-interface	324
	destination-port	325
	fields (for Interface Profiles)	326
	fields (for Routing Engine Profiles)	327
	file (Associating with a Profile)	328
	file (Configuring a Log File)	329
	files	330
	filter-profile	331
	flow (Security Flow)	332
	global-threshold	334
	global-weight	335
	hardware-timestamp	335
	icmp	336
	idp (Security Alarms)	336
	inet6-options (Services)	337
	interface-profile	338
	interval	339
	ip-monitoring	340
	ip-monitoring (Services)	341
	maximum-capture-size (Datapath Debug)	342
	mib-profile	343
	mpls (Security Forwarding Options)	344

	next-hop	344
	nonpersistent	345
	object-names	345
	operation	346
	packet-capture	347
	packet-filter	348
	probe	349
	probe-interval	350
	probe-limit	351
	probe-server	352
	probe-type	353
	redundancy-group (Chassis Cluster)	354
	retry-interval (Chassis Cluster)	355
	routing-engine-profile	356
	rpm (Services)	357
	Security Configuration Statement Hierarchy	358
	size	360
	source-classes	360
	start-time	361
	target (Services RPM)	362
	thresholds	363
	traceoptions (Security Datapath Debug)	365
	transfer-interval	366
	traps	367
Chapter 21	Operational Commands	369
	clear chassis cluster ip-monitoring failure-count	371
	clear chassis cluster ip-monitoring failure-count ip-address	372
	monitor list	373
	monitor start	374
	monitor stop	376
	mtrace monitor	377
	ping mpls l2circuit	379
	ping mpls l2vpn	382
	ping mpls l3vpn	385
	ping mpls ldp	388
	ping mpls lsp-end-point	391
	ping mpls rsvp	393
	request pppoe connect	398
	request pppoe disconnect	399
	request services ip-monitoring preempt-restore policy	400
	show chassis alarms	401
	show configuration	403
	show chassis cluster ip-monitoring status redundancy-group	406
	show interfaces (SRX Series)	409
	show poe interface (View)	440
	show poe telemetries	442
	show pppoe interfaces	444
	show pppoe statistics	448

show security alarms	450
show security datapath-debug capture	454
show security datapath-debug counter	455
show security monitoring	456
show security monitoring fpc fpc-number	458
show security monitoring performance session	461
show security monitoring performance spu	462
show services ip-monitoring status	464
show services rpm probe-results (View)	468
show system alarms	473
traceroute	474

Part 6

Index

Index	481
-----------------	-----

List of Figures

Part 2	Configuring Monitoring Options	
Chapter 5	Using RPM to Measure Network Performance	49
	Figure 1: Sample RPM Graphs	73
Chapter 6	Configuring IP Monitoring	77
	Figure 2: IP Monitoring on a High-End SRX Series Device Topology Example . . .	83
Part 4	Troubleshooting	
Chapter 19	Troubleshooting Security Devices	291
	Figure 3: PPP and MLPPP Headers	298

List of Tables

	About the Documentation	xix
	Table 1: Notice Icons	xxi
	Table 2: Text and Syntax Conventions	xxi
Part 1	Overview	
Chapter 1	Introduction to Network Monitoring	3
	Table 3: J-Web Interface Troubleshoot Options	5
	Table 4: CLI Diagnostic Command Summary	5
Chapter 2	Accounting Options, Source Class Usage, and Destination Class Usage Overview	7
	Table 5: Types of Accounting Profiles	7
Part 2	Configuring Monitoring Options	
Chapter 4	Configuring Interface Alarms	37
	Table 6: Interface Alarm Conditions	39
	Table 7: System Alarm Conditions and Corrective Actions	42
	Table 8: Alarms Monitoring Page	47
Chapter 5	Using RPM to Measure Network Performance	49
	Table 9: RPM Statistics	51
	Table 10: RPM Configuration Summary	68
	Table 11: Summary of Key RPM Output Fields	73
Chapter 6	Configuring IP Monitoring	77
	Table 12: Test Parameters and Default Values	78
	Table 13: Threshold Supported and Description	79
Part 3	Monitoring Common Security Features	
Chapter 7	Displaying Real-Time Information from Device to Host	95
	Table 14: CLI mtrace from-source Command Options	95
	Table 15: CLI mtrace from-source Command Output Summary	97
	Table 16: CLI traceroute monitor Command Options	98
	Table 17: CLI traceroute monitor Command Output Summary	99
Chapter 8	Monitoring Application Layer Gateways Features	101
	Table 18: Summary of Key H.323 Counters Output Fields	101
	Table 19: Summary of Key MGCP Calls Output Fields	103
	Table 20: Summary of Key MGCP Counters Output Fields	104
	Table 21: Summary of Key MGCP Endpoints Output Fields	105

	Table 22: Summary of Key SCCP Calls Output Fields	106
	Table 23: Summary of Key SCCP Counters Output Fields	106
	Table 24: Summary of Key SIP Calls Output Fields	108
	Table 25: Summary of Key SIP Counters Output Fields	109
	Table 26: Summary of Key SIP Rate Output Fields	111
	Table 27: Summary of Key SIP Transactions Output Fields	112
	Table 28: ALG H.323 Monitoring Page	112
	Table 29: Voice ALG MGCP Monitoring Page	115
	Table 30: Voice ALG SCCP Monitoring Page	117
	Table 31: Voice ALG SIP Monitoring Page	120
	Table 32: Voice ALG Summary Monitoring Page	125
Chapter 9	Monitoring Class of Service	127
	Table 33: Summary of Key CoS Interfaces Output Fields	128
	Table 34: Summary of Key CoS Classifier Output Fields	128
	Table 35: Summary of Key CoS Value Alias Output Fields	130
	Table 36: Summary of Key CoS RED Drop Profile Output Fields	130
	Table 37: Summary of Key CoS Forwarding Class Output Fields	132
	Table 38: Summary of Key CoS Rewrite Rules Output Fields	132
	Table 39: Summary of Key CoS Scheduler Maps Output Fields	133
	Table 40: Summary of Key CoS Classifier Output Fields	135
Chapter 10	Monitoring Interfaces and Switching Functions	137
	Table 41: CLI monitor interface Output Control Keys	138
	Table 42: CLI monitor interface traffic Output Control Keys	138
	Table 43: Address Pools Monitoring Page	139
	Table 44: Summary of Ethernet Switching Output Fields	141
	Table 45: GVRP Monitoring Page	142
	Table 46: Summary of Key MPLS Interface Information Output Fields	144
	Table 47: Summary of Key MPLS LSP Information Output Fields	144
	Table 48: Summary of Key MPLS LSP Statistics Output Fields	146
	Table 49: Summary of Key RSVP Session Information Output Fields	147
	Table 50: Summary of Key RSVP Interfaces Information Output Fields	148
	Table 51: Summary of Key PPPoE Output Fields	150
	Table 52: Spanning Tree Monitoring Page	153
Chapter 11	Monitoring NAT	155
	Table 53: Source NAT Monitoring Page	155
	Table 54: Summary of Key Destination NAT Output Fields	161
	Table 55: Summary of Key Static NAT Output Fields	163
	Table 56: Summary of Key Incoming Table Output Fields	165
	Table 57: Summary of Key Interface NAT Output Fields	165
Chapter 12	Monitoring Security Policies	167
	Table 58: Filtering Route Messages	169
	Table 59: Summary of Key Routing Information Output Fields	169
	Table 60: Summary of Key RIP Routing Output Fields	170
	Table 61: Summary of Key OSPF Routing Output Fields	172
	Table 62: Summary of Key BGP Routing Output Fields	174
	Table 63: View Policy Log Fields	175

	Table 64: Policy Events Detail Fields	177
	Table 65: Security Policies Monitoring Output Fields	178
	Table 66: Check Policies Output	181
	Table 67: Summary of Key Screen Counters Output Fields	183
	Table 68: Summary of IDP Status Output Fields	186
	Table 69: Summary of Key Flow Gate Output Fields	187
	Table 70: Summary of Key Firewall Authentication Table Output Fields	188
	Table 71: Summary of Key Firewall Authentication History Output Fields	189
	Table 72: Summary of Dot1X Output Fields	191
Chapter 13	Monitoring Events, Services and System	193
	Table 73: Summary of Key DHCP Client Binding Output Fields	193
	Table 74: Events Monitoring Page	194
Chapter 14	Monitoring Unified Threat Management Features	203
	Table 75: Statistics Tab Output in the Threats Report	208
	Table 76: Activities Tab Output in the Threats Report	210
	Table 77: Traffic Report Output	213
Chapter 15	Monitoring VPNs	217
	Table 78: Summary of Key IKE SA Information Output Fields	217
	Table 79: IPsec VPN—Phase I Monitoring Page	221
	Table 80: IPsec VPN—Phase II Monitoring Page	222
	Table 81: Summary of Key IPsec VPN Information Output Fields	223
Part 4	Troubleshooting	
Chapter 16	Configuring Data Path Debugging and Trace Options	231
	Table 82: CLI mtrace monitor Command Output Summary	240
	Table 83: Traceroute Field Summary	241
	Table 84: J-Web Traceroute Results and Output Summary	243
	Table 85: CLI traceroute Command Options	245
Chapter 17	Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits	247
	Table 86: Options for Checking MPLS Connections	248
	Table 87: CLI ping Command Options	250
	Table 88: J-Web Ping Host Field Summary	252
	Table 89: Ping Host Results and Output	254
	Table 90: J-Web Ping MPLS Field Summary	255
	Table 91: J-Web Ping MPLS Results and Output Summary	258
	Table 92: CLI ping mpls l2circuit Command Options	259
	Table 93: CLI ping mpls l2vpn Command Options	260
	Table 94: CLI ping mpls l3vpn Command Options	261
	Table 95: CLI ping mpls ldp and ping mpls lsp-end-point Command Options	262
Chapter 18	Using Packet Capture to Analyze Network Traffic	265
	Table 96: CLI monitor traffic Command Options	281
	Table 97: CLI monitor traffic Match Conditions	283
	Table 98: CLI monitor traffic Logical Operators	284
	Table 99: CLI monitor traffic Arithmetic, Binary, and Relational Operators	284
	Table 100: Packet Capture Field Summary	286

	Table 101: J-Web Packet Capture Results and Output Summary	288
Chapter 19	Troubleshooting Security Devices	291
	Table 102: CoS Components Applied on Multilink Bundles and Constituent Links	294
	Table 103: PPP and MLPPP Encapsulation Overhead	298
	Table 104: Number of Packets Transmitted on a Queue	301
Part 5	Configuration Statements and Operational Commands	
Chapter 21	Operational Commands	369
	Table 105: monitor list Output Fields	373
	Table 106: monitor start Output Fields	374
	Table 107: mtrace monitor Output Fields	377
	Table 108: show chassis alarms Output Fields	401
	Table 109: show chassis cluster ip-monitoring status Output Fields	406
	Table 110: show chassis cluster ip-monitoring status redundancy group Reason Fields	407
	Table 111: show interfaces Output Fields	412
	Table 112: show poe interface Output Fields	440
	Table 113: show poe telemetries interface Output Fields	442
	Table 114: show pppoe interfaces Output Fields	444
	Table 115: show pppoe statistics Output Fields	448
	Table 116: show security alarms	451
	Table 117: show security monitoring fpc fpc-number Output Fields	458
	Table 118: show services ip-monitoring status Output Fields	464
	Table 119: show services rpm probe-results Output Fields	468
	Table 120: traceroute Output Fields	476

About the Documentation

- Documentation and Release Notes on page xix
- Supported Platforms on page xix
- Using the Examples in This Manual on page xix
- Documentation Conventions on page xxi
- Documentation Feedback on page xxiii
- Requesting Technical Support on page xxiii

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:

- SRX Series
- vSRX

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 xxi](#) 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 xxi](#) 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

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

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<pre>user@host> show chassis alarms</pre> <p>No alarms currently active</p>
<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.	<p>Configure the machine's domain name:</p> <pre>[edit] root@# set system domain-name domain-name</pre>
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 <code>[edit protocols ospf area area-id]</code> 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 string2 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 [community-ids]
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.

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

Convention	Description	Examples
> (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: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- 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://tools.juniper.net/SerialNumberEntitlementSearch/>

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

- [Introduction to Network Monitoring on page 3](#)
- [Accounting Options, Source Class Usage, and Destination Class Usage Overview on page 7](#)
- [Gathering Statistics for Accounting Purposes on page 11](#)

CHAPTER 1

Introduction to Network Monitoring

- [Monitoring Overview on page 3](#)
- [Diagnostic Tools Overview on page 4](#)

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
```

<code>display</code>	Show additional kinds of information
<code>except</code>	Show only text that does not match a pattern
<code>find</code>	Search for first occurrence of pattern
<code>hold</code>	Hold text without exiting the prompt
<code>last</code>	Display end of output only
<code>match</code>	Show only text that matches a pattern
<code>no-more</code>	Don't paginate output
<code>request</code>	Make system-level requests
<code>resolve</code>	Resolve IP addresses
<code>save</code>	Save output text to file
<code>trim</code>	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 142](#)
 - [Diagnostic Tools Overview on page 4](#)

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 4](#)
- [CLI Diagnostic Commands on page 5](#)

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 3 on page 5](#) describes the functions of the Troubleshoot options.

Table 3: 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 4 on page 5](#).

Table 4: 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.

Table 4: CLI Diagnostic Command Summary (*continued*)

Command	Function
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.
traceroute	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 247](#)
 - [Configuring Ping MPLS on page 249](#)
 - [Using the J-Web Ping Host Tool on page 252](#)
 - [Using the ping Command on page 250](#)

CHAPTER 2

Accounting Options, Source Class Usage, and Destination Class Usage Overview

- [Accounting Options Overview on page 7](#)
- [Understanding Source Class Usage and Destination Class Usage Options on page 8](#)

Accounting Options Overview

Supported Platforms [ACX Series](#), [M Series](#), [MX 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 5 on page 7](#).

Table 5: 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](#)

- [Accounting Options Configuration on page 11](#)
- [Configuring Accounting-Data Log Files on page 14](#)
- [Configuring the Interface Profile on page 17](#)
- [Configuring the Filter Profile on page 20](#)
- *Configuration Statements at the [edit accounting-options] Hierarchy Level*

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 only when third-generation FPCs are installed on the router and *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 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 for Routing Devices*, the *Junos OS Network Interfaces Library for Routing Devices*, and the *Junos OS, Release 15.1*.

**Related
Documentation**

- [Example: Grouping Source and Destination Prefixes into a Forwarding Class](#)
- [Configuring SCU or DCU on page 24](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 26](#)
- [Configuring Class Usage Profiles on page 28](#)
- [Configuring the MIB Profile on page 30](#)
- [Configuring the Routing Engine Profile on page 32](#)

CHAPTER 3

Gathering Statistics for Accounting Purposes

- [Accounting Options Configuration on page 11](#)
- [Configuring Accounting-Data Log Files on page 14](#)
- [Configuring the Interface Profile on page 17](#)
- [Configuring the Filter Profile on page 20](#)
- [Example: Configuring a Filter Profile on page 22](#)
- [Example: Configuring Interface-Specific Firewall Counters and Filter Profiles on page 23](#)
- [Configuring SCU or DCU on page 24](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 26](#)
- [Configuring Class Usage Profiles on page 28](#)
- [Configuring the MIB Profile on page 30](#)
- [Configuring the Routing Engine Profile on page 32](#)

Accounting Options Configuration

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

This topic contains the following sections:

- [Accounting Options—Full Configuration on page 11](#)
- [Minimum Accounting Options Configuration on page 13](#)

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;
  }
}
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 it 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;
  }
  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 7](#)
- [Understanding Device Management Functions in Junos OS](#)
- [Configuring Accounting-Data Log Files on page 14](#)
- [Configuring the Interface Profile on page 17](#)
- [Configuring the Filter Profile on page 20](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level](#)

Configuring Accounting-Data Log Files

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

An accounting profile specifies what statistics should be collected and written 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]
file filename {
  archive-sites {
    site-name;
  }
  files number;
  nonpersistent;
  size bytes;
  start-time time;
  transfer-interval minutes;
}
```

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 the Storage Location of the File on page 15](#)
- [Configuring the Maximum Size of the File on page 16](#)
- [Configuring the Maximum Number of Files on page 16](#)
- [Configuring the Start Time for File Transfer on page 16](#)
- [Configuring the Transfer Interval of the File on page 16](#)
- [Configuring Archive Sites on page 17](#)

Configuring the Storage Location of the File

To configure the storage location of the files in the **mfs/var/log** directory (on DRAM) instead of the **cf/var/log** directory (on the compact flash drive), include the **nonpersistent** statement at the **[edit accounting-options file filename]** hierarchy level:

```
[edit accounting-options file filename]
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. Therefore, you should back up these files periodically.

Configuring the Maximum Size of the File

To configure the maximum size of the files, include the **size** statement at the **[edit accounting-options file *filename*]** hierarchy level:

```
[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 the Maximum Number of Files

To configure the maximum number of files, include the **files** statement at the **[edit accounting-options file *filename*]** hierarchy level:

```
[edit accounting-options file filename]  
files number;
```

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 **number** is 3 and the default value is 10.

Configuring the Start Time for File Transfer

To configure the start time for transferring files, include the **start-time** statement at the **[edit accounting-options file *filename*]** hierarchy level:

```
[edit accounting-options file filename]  
start-time time;
```

The start-time statement specifies a start time for file transfer (**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 transfer interval of the files, include the **transfer-interval** statement at the **[edit accounting-options file *filename*]** hierarchy level:

```
[edit accounting-options file filename]  
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 or not
- Whether an archive site is configured or not

When you have a relatively smaller 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.

Configuring Archive Sites

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 archive sites, include the **archive-sites** statement at the **[edit accounting-options file filename]** hierarchy level:

```
[edit accounting-options file filename]
archive-sites {
  site-name;
}
```

site-name is any valid FTP URL. For more information about specifying valid FTP URLs, see the *Junos OS Administration Library for Routing Devices*. 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**.

Related Documentation

- [Accounting Options Overview on page 7](#)
- [Understanding Device Management Functions in Junos OS](#)
- [Accounting Options Configuration on page 11](#)
- [Configuring the Interface Profile on page 17](#)
- [Configuring the Filter Profile on page 20](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level](#)

Configuring the Interface Profile

Supported Platforms **M Series, MX 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 for Routing Devices*.

To configure an interface profile, perform the tasks described in the following sections:

- [Configuring Fields on page 18](#)
- [Configuring the File Information on page 18](#)
- [Configuring the Interval on page 19](#)
- [Example: Configuring the Interface Profile on page 19](#)

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 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,xen-1/0/0.0,8,134696815,3681534,501088,40723,0,0
if_profile1,976823538,xen-1/0/0.7,134696815,3681534,501088
...
#FILE CLOSED 976824378 2000-12-14-20:06:18

```

Related Documentation

- [Accounting Options Overview on page 7](#)
- [Understanding Device Management Functions in Junos OS](#)
- [Accounting Options Configuration on page 11](#)
- [Configuring Accounting-Data Log Files on page 14](#)
- [Configuring the Filter Profile on page 20](#)
- [Configuration Statements at the \[edit accounting-options\] Hierarchy Level](#)

Configuring the Filter Profile

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

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 21](#)
- [Configuring the File Information on page 21](#)
- [Configuring the Interval on page 21](#)

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 7](#)
- [Understanding Device Management Functions in Junos OS](#)
- [Accounting Options Configuration on page 11](#)
- [Configuring Accounting-Data Log Files on page 14](#)

Example: Configuring a Filter Profile

Supported Platforms M Series, MX 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 20](#)
- [Example: Configuring Interface-Specific Firewall Counters and Filter Profiles on page 23](#)

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 20](#)
- [Configuring the Interface Profile on page 17](#)

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 24](#)
- [Applying the Policy to the Forwarding Table on page 25](#)
- [Enabling Accounting on Inbound and Outbound Interfaces on page 25](#)

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 for Routing Devices* and the *Junos OS Network Interfaces Library for Routing Devices*.

Related Documentation

- [Understanding Source Class Usage and Destination Class Usage Options on page 8](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 26](#)
- [Configuring Class Usage Profiles on page 28](#)
- [Configuring the MIB Profile on page 30](#)
- [Configuring the Routing Engine Profile on page 32](#)

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 26](#)
- [Example: Mapping the VRF Instance Type to the Virtual Loopback Tunnel Interface on page 27](#)
- [Example: Sending Traffic Received from the Virtual Loopback Interface Out the Source Class Output Interface on page 27](#)

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 8](#)
- [Configuring SCU or DCU on page 24](#)

- [Configuring Class Usage Profiles on page 28](#)
- [Configuring the MIB Profile on page 30](#)
- [Configuring the Routing Engine Profile on page 32](#)

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 28](#)
- [Configuring the File Information on page 28](#)
- [Configuring the Interval on page 29](#)
- [Creating a Class Usage Profile to Collect Source Class Usage Statistics on page 29](#)
- [Creating a Class Usage Profile to Collect Destination Class Usage Statistics on page 29](#)

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-profile1;  
  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 8](#)
- [Configuring SCU or DCU on page 24](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 26](#)
- [Configuring the Routing Engine Profile on page 32](#)

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 30](#)
- [Configuring the Interval on page 31](#)
- [Configuring the MIB Operation on page 31](#)
- [Configuring MIB Object Names on page 31](#)
- [Example: Configuring a MIB Profile on page 31](#)

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: 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;
}
}
```

Related Documentation

- [Understanding Source Class Usage and Destination Class Usage Options on page 8](#)
- [Configuring SCU or DCU on page 24](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 26](#)
- [Configuring Class Usage Profiles on page 28](#)
- [Configuring the Routing Engine Profile on page 32](#)

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 32](#)
- [Configuring the File Information on page 33](#)
- [Configuring the Interval on page 33](#)
- [Example: Configuring a Routing Engine Profile on page 33](#)

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 8](#)
- [Configuring SCU or DCU on page 24](#)
- [Configuring SCU on a Virtual Loopback Tunnel Interface on page 26](#)
- [Configuring Class Usage Profiles on page 28](#)
- [Configuring the MIB Profile on page 30](#)

PART 2

Configuring Monitoring Options

- [Configuring Interface Alarms on page 37](#)
- [Using RPM to Measure Network Performance on page 49](#)
- [Configuring IP Monitoring on page 77](#)

CHAPTER 4

Configuring Interface Alarms

- [Alarm Overview on page 37](#)
- [Example: Configuring Interface Alarms on page 43](#)
- [Monitoring Active Alarms on a Device on page 46](#)
- [Monitoring Alarms on page 47](#)

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 37](#)
- [Alarm Severity on page 38](#)
- [Alarm Conditions on page 38](#)

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, 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 39](#)
- [System Alarm Conditions on page 42](#)

Interface Alarm Conditions

[Table 6 on page 39](#) 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 6: 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 6: 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 6: 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 6: 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 7 on page 42 lists the two preset system alarms, the condition that triggers each alarm, and the action you take to correct the condition.

Table 7: System Alarm Conditions and Corrective Actions

Alarm Type	Alarm Condition	Corrective Action
Configuration	The rescue configuration is not set.	Set the rescue configuration.

Table 7: 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.

- Related Documentation**
- [Example: Configuring Interface Alarms on page 43](#)
 - [Monitoring Active Alarms on a Device on page 46](#)
 - [Monitoring Alarms on page 47](#)
 - [System Log Messages](#)

Example: Configuring Interface Alarms

Supported Platforms [SRX Series](#)

This example shows how to configure interface alarms.

- [Requirements on page 43](#)
- [Overview on page 43](#)
- [Configuration on page 44](#)
- [Verification on page 45](#)

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 37](#).

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

- [Alarm Overview on page 37](#)
- [Monitoring Active Alarms on a Device on page 46](#)
- [Monitoring Alarms on page 47](#)

Monitoring Active Alarms on a Device

Supported Platforms [SRX Series, vSRX](#)

Purpose Use to monitor and filter alarms on a Juniper Networks device.

Action 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:

- 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.

To filter which alarms appear, use the following options:

- 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 37](#)
- [Example: Configuring Interface Alarms on page 43](#)
- [Monitoring Alarms on page 47](#)

Monitoring Alarms

Supported Platforms SRX Series, vSRX

Purpose Use the monitoring functionality to view the alarms page.

Action To monitor alarms select **Monitor>Events and Alarms>View Alarms** in the J-Web user interface.

Meaning Table 8 on page 47 summarizes key output fields in the alarms page.

Table 8: 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. 	—
Severity	Specifies the alarm severity that you want to monitor <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.	—

Table 8: Alarms Monitoring Page (continued)

Field	Value	Additional Information
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 46](#)
 - [Monitoring Events on page 193](#)
 - [Monitoring Security Events by Policy on page 175](#)

CHAPTER 5

Using RPM to Measure Network Performance

- [RPM Overview on page 49](#)
- [IPv6 RPM Probes on page 53](#)
- [Guidelines for Configuring RPM Probes for IPv6 on page 54](#)
- [RPM Support for VPN Routing and Forwarding on page 55](#)
- [Example: Configuring Basic RPM Probes on page 55](#)
- [Example: Configuring RPM Using TCP and UDP Probes on page 59](#)
- [Example: Configuring RPM Probes for BGP Monitoring on page 62](#)
- [Directing RPM Probes to Select BGP Devices on page 65](#)
- [Configuring RPM Timestamping on page 65](#)
- [Configuring IPv6 RPM Probes on page 66](#)
- [Tuning RPM Probes on page 67](#)
- [RPM Configuration Options on page 68](#)
- [Monitoring RPM Probes on page 72](#)

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 50](#)
- [RPM Tests on page 50](#)
- [Probe and Test Intervals on page 50](#)
- [Jitter Measurement with Hardware Timestamping on page 51](#)
- [RPM Statistics on page 51](#)
- [RPM Thresholds and Traps on page 53](#)
- [RPM for BGP Monitoring on page 53](#)

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.

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 Low Memory devices and SRX340 High Memory devices, the RPM server operation does not work when the probe is configured with the option destination-interface.

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 9 on page 51](#).

Table 9: RPM Statistics

RPM Statistics	Description
Round-Trip Times	

Table 9: RPM Statistics (*continued*)

RPM Statistics	Description
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
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 68](#)
- [RPM Support for VPN Routing and Forwarding on page 55](#)
- [Example: Configuring Basic RPM Probes on page 55](#)
- [Monitoring RPM Probes on page 72](#)
- [Determine What Causes Jitter and Latency on the Multilink Bundle on page 295](#)

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 54](#)
- [Configuring IPv6 RPM Probes on page 66](#)

Guidelines for Configuring RPM Probes for IPv6

Supported Platforms **vSRX**

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
6to4	2002::/16
6bone	5f00::/8
ORCHID	2001:10::/28
Teredo	2001::/32

IPv6 Address Type	IPv6 Address Prefix
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 66](#)

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 49](#)
- [RPM Configuration Options on page 68](#)
- [Monitoring RPM Probes on page 72](#)

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 56](#)
- [Overview on page 56](#)

- [Configuration on page 56](#)
- [Verification on page 58](#)

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.

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 58](#)
- [Verifying RPM Statistics on page 58](#)

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 49](#)
 - [RPM Configuration Options on page 68](#)
 - [Tuning RPM Probes on page 67](#)

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 59](#)
- [Overview on page 60](#)
- [Configuration on page 60](#)
- [Verification on page 62](#)

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 55](#).

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-0/0/0.0
```
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.

<pre>user@host> show services rpm active-servers
Protocol: TCP, Port: 50000

Protocol: UDP, Port: 50037</pre> |
| Related Documentation | <ul style="list-style-type: none">• RPM Overview on page 49• RPM Configuration Options on page 68• Example: Configuring Basic RPM Probes on page 55• Example: Configuring RPM Probes for BGP Monitoring on page 62• Tuning RPM Probes on page 67 |

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 62](#)
- [Overview on page 63](#)
- [Configuration on page 63](#)
- [Verification on page 64](#)

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 55](#).
- 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 59](#).

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 49](#)
 - [RPM Configuration Options on page 68](#)
 - [Directing RPM Probes to Select BGP Devices on page 65](#)
 - [Tuning RPM Probes on page 67](#)

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 49](#)
 - [RPM Configuration Options on page 68](#)
 - [Example: Configuring Basic RPM Probes on page 55](#)
 - [Example: Configuring RPM Probes for BGP Monitoring on page 62](#)
 - [Tuning RPM Probes on page 67](#)

Configuring RPM Timestamping

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You can timestamp the following RPM probe types: **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp**.

This example shows how to enable timestamping for customerA. The test for customerA is identified as customerA-test.

To configure timestamping:

1. Specify the RPM probe owner for which you want to enable timestamping.

```
[edit services rpm]
user@host# edit probe customerA
```

2. Specify a name for the test.

```
[edit services rpm probe customerA]
user@host# edit test customerA-test
```

3. Enable timestamping.

```
[edit services rpm probe customerA test customerA-test]
user@host# edit hardware-timestamp
```

4. (Optional) If preferred, indicate that you want timestamping to be only one-way.

```
[edit services rpm probe customerA test customerA-test]
user@host# edit one-way-hardware-timestamp
```



NOTE: You cannot include both the `source-address` and `hardware-timestamp` or `one-way-hardware-timestamp` statements at the `[edit services rpm probe probe-name test test-name]` hierarchy level simultaneously.

Related Documentation

- [RPM Overview on page 49](#)
- [RPM Configuration Options on page 68](#)
- [Example: Configuring Basic RPM Probes on page 55](#)
- [Example: Configuring RPM Using TCP and UDP Probes on page 59](#)
- [Tuning RPM Probes on page 67](#)

Configuring IPv6 RPM Probes

Supported Platforms **vSRX**

You can configure IPv6 source and 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 source address for the test.

```
[edit services rpm probe customerA test ipv6-test]
user@host# set source-address 2001:db8:1a:1112::20
```

5. Specify the target address for the test.

```
[edit services rpm probe customerA test ipv6-test]
user@host# set target inet6-address 2001:db8:1a:1112::1
```

6. Configure the remaining RPM test parameters.

Related Documentation

- [Guidelines for Configuring RPM Probes for IPv6 on page 54](#)

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 55](#).

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 49](#)
 - [RPM Configuration Options on page 68](#)
 - [Example: Configuring RPM Probes for BGP Monitoring on page 62](#)

RPM Configuration Options

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

You can configure real-time performance monitoring (RPM) parameters. See [Table 10 on page 68](#) for a summary of the configuration options.

Table 10: 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.
Request Information		

Table 10: RPM Configuration Summary (*continued*)

Field	Function	Your Action
Probe Type (required)	Specifies the type of probe to send as part of the test.	<p>Select the desired probe type from the list:</p> <ul style="list-style-type: none"> • http-get • http-get-metadata • icmp6-ping • icmp-ping • icmp-ping-timestamp • tcp-ping • udp-ping
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	<p>Specifies the TCP or UDP port to which probes are sent.</p> <p>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.</p>	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	<p>Enables timestamping of RPM probe messages. You can timestamp the following RPM probes to improve the measurement of latency or jitter:</p> <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.

Maximum Probe Thresholds

Table 10: RPM Configuration Summary (*continued*)

Field	Function	Your Action
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		
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.

Table 10: RPM Configuration Summary (*continued*)

Field	Function	Your Action
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		
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.

Table 10: RPM Configuration Summary (*continued*)

Field	Function	Your Action
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.

Related Documentation

- [RPM Overview on page 49](#)
- [Example: Configuring Basic RPM Probes on page 55](#)
- [Example: Configuring RPM Using TCP and UDP Probes on page 59](#)
- [Example: Configuring RPM Probes for BGP Monitoring on page 62](#)

Monitoring RPM Probes

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

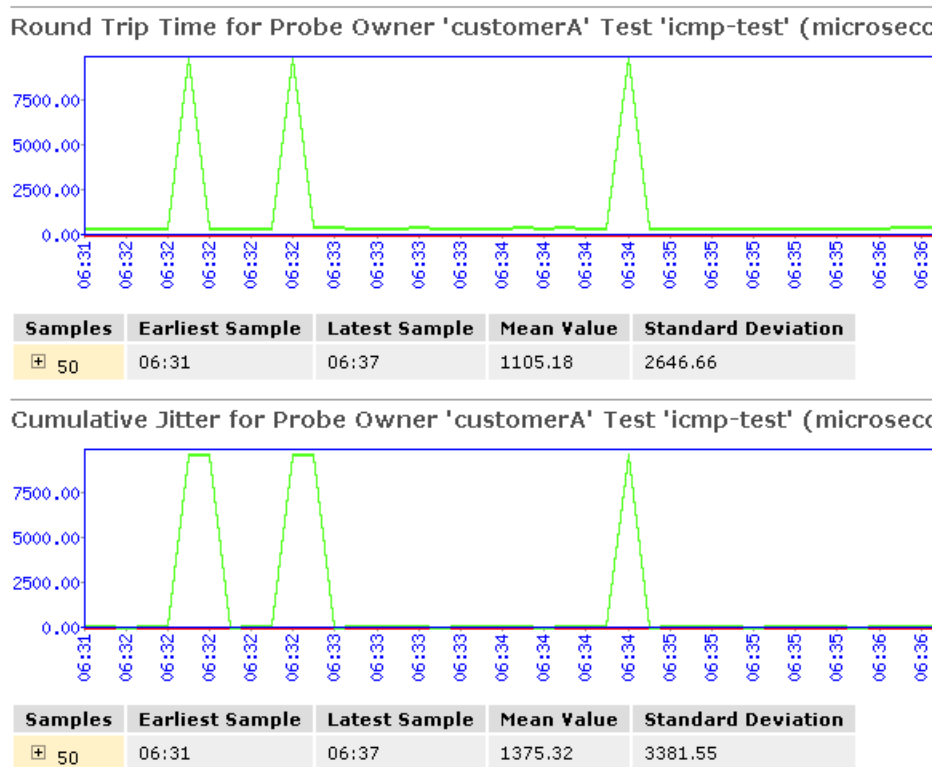
The RPM information includes the round-trip time, jitter, and standard deviation values for each configured RPM test on the device. To view these RPM properties, select **Troubleshoot > RPM > View RPM** in the J-Web user interface, or in configuration mode enter the **show** command:

[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 1 on page 73](#) shows sample graphs for an RPM test.

Figure 1: Sample RPM Graphs



In [Figure 1 on page 73](#), 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 11 on page 73](#) summarizes key output fields in RPM displays.

Table 11: 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 11: 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 11: 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 49](#)
 - [RPM Support for VPN Routing and Forwarding on page 55](#)
 - [RPM Configuration Options on page 68](#)

CHAPTER 6

Configuring IP Monitoring

- [IP Monitoring Overview on page 77](#)
- [Understanding IP Monitoring Test Parameters on page 78](#)
- [Example: Configuring IP Monitoring on Branch SRX Series Devices on page 79](#)
- [Understanding IP Monitoring Through Redundant Ethernet Interface Link Aggregation Groups on page 81](#)
- [Example: Configuring IP Monitoring on High-End SRX Series Devices on page 82](#)
- [Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring on page 87](#)

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 78](#)

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 12 on page 78](#) lists the test parameters and its default values:

Table 12: Test Parameters and Default Values

Parameter	Default Value
probe-count	1
probe-interval	3 seconds
test-interval	1 second

Table 13 on page 79 lists the supported threshold and its description:

Table 13: 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 77](#)

Example: Configuring IP Monitoring on Branch SRX Series Devices

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

This example shows how to monitor IP on branch SRX Series devices.

- [Requirements on page 79](#)
- [Overview on page 79](#)
- [Configuration on page 79](#)
- [Verification on page 81](#)

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 for the branch 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 77](#)
- [Understanding IP Monitoring Test Parameters on page 78](#)

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 77](#)

Example: Configuring IP Monitoring on High-End SRX Series Devices

Supported Platforms [SRX1500, SRX5600, SRX5800, vSRX](#)

This example shows how to monitor IP on a high-end SRX Series device with chassis cluster enabled.

- [Requirements on page 82](#)
- [Overview on page 82](#)
- [Configuration on page 83](#)
- [Verification on page 85](#)

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 a high-end SRX Series device.




NOTE: IP monitoring is not supported on an NP-IOC card.

Topology

[Figure 2 on page 83](#) shows the topology used in this example.

Figure 2: IP Monitoring on a High-End SRX Series Device Topology Example



In this example, two SRX5800 devices in a chassis cluster are connected to an SRX650 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 a High-End SRX Series Device on page 84](#)

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 111.111.111.111
weight 80
set chassis cluster redundancy-group 1 ip-monitoring family inet 111.111.111.111
interface reth0.0 secondary-ip-address 12.1.1.3
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 12.1.1.1/24
set routing-options static route 111.111.111.111/32 next-hop 12.1.1.2
```

Configuring IP Monitoring on a High-End 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 a high-end 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 12.1.1.1/24
```
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 111.111.111.111/32 next-hop 12.1.1.2
```
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  
111.111.111.111 weight 80
```

```
user@host# set chassis cluster redundancy-group 1 ip-monitoring family inet  
111.111.111.111 interface reth0.0 secondary-ip-address 12.1.1.3
```



NOTE:

- The redundant Ethernet (reth0) IP address, 12.1.1.1/24, is used to send ICMP packets from node 0 to check the reachability of the monitored IP.
- The secondary IP address, 12.1.1.3, 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 85](#)
- [Verifying Chassis Cluster IP Monitoring Status—Before Failover on page 86](#)
- [Verifying Chassis Cluster Status—After Failover on page 86](#)
- [Verifying Chassis Cluster IP Monitoring Status—After Failover on page 87](#)

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
111.111.111.111	reachable	0	n/a

node1:

Redundancy group: 1

IP address	Status	Failure count	Reason
111.111.111.111	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
111.111.111.111	unreachable	1	unknown

```
node1:
```

```
Redundancy group: 1
```

IP address	Status	Failure count	Reason
111.111.111.111	reachable	0	n/a

Related Documentation

- [Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway](#)

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 88](#)
- [Overview on page 88](#)

- [Configuration on page 89](#)
- [Verification on page 90](#)

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 Branch SRX Series Devices](#) or [Example: Setting the Chassis Cluster Node ID and Cluster ID](#).
- 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, and 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 for Branch SRX Series Devices](#)
 - [Understanding Chassis Cluster Redundancy Group IP Address Monitoring for High-End SRX Series Devices](#)

- *Understanding Chassis Cluster Redundancy Group Failover*

PART 3

Monitoring Common Security Features

- [Displaying Real-Time Information from Device to Host on page 95](#)
- [Monitoring Application Layer Gateways Features on page 101](#)
- [Monitoring Class of Service on page 127](#)
- [Monitoring Interfaces and Switching Functions on page 137](#)
- [Monitoring NAT on page 155](#)
- [Monitoring Security Policies on page 167](#)
- [Monitoring Events, Services and System on page 193](#)
- [Monitoring Unified Threat Management Features on page 203](#)
- [Monitoring VPNs on page 217](#)

CHAPTER 7

Displaying Real-Time Information from Device to Host

- [Displaying Multicast Path Information on page 95](#)
- [Displaying Real-Time Monitoring Information on page 97](#)

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> <ttl number> <wait-time seconds> <loop>
<multicast-response | unicast-response> <no-resolve> <no-router-alert> <brief |
detail>
```

[Table 14 on page 95](#) describes the **mtrace from-source** command options.

Table 14: CLI mtrace from-source Command Options

Option	Description
source host	Traces the path to the specified hostname or IP address.
extra-hops number	(Optional) Sets the number of extra hops to trace past nonresponsive devices. Specify a value from 0 through 255.
group address	(Optional) Traces the path for the specified group address. The default value is 192.0.2.0.
interval seconds	(Optional) Sets the interval between statistics gathering. The default value is 10.
max-hops number	(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 number	(Optional) Sets the maximum number of query attempts for any hop. Specify a value from 1 through 32. The default value is 3.

Table 14: CLI mtrace from-source Command Options (*continued*)

Option	Description
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 Packet
192.1.4.1 To 224.1.1.1 v ___/ rtt 16 ms Rate Lost/Sent =
Pct Rate 192.168.195.37 192.1.40.2 routerC.mycompany.net v ^
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 15 on page 97 summarizes the output fields of the display.



NOTE: The packet statistics gathered from Juniper Networks devices and routing nodes always display as 0.

Table 15: 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 milliseconds ms	Total time between the sending of the query packet and the receiving of the response packet.
total ttl of number 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 3](#)

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 16 on page 98 describes the **traceroute monitor** command options.

Table 16: 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.
<i>summary</i>	(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                                     Loss%   Snt
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 17 on page 99 summarizes the output fields of the display.

Table 17: 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.
quit	Quits the traceroute monitor command. Press Q to quit the traceroute monitor command.
Packets	
number	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 239](#)

CHAPTER 8

Monitoring Application Layer Gateways Features

- [Monitoring H.323 ALG Information on page 101](#)
- [Monitoring MGCP ALGs on page 102](#)
- [Monitoring SCCP ALGs on page 105](#)
- [Monitoring SIP ALGs on page 108](#)
- [Monitoring Voice ALG H.323 on page 112](#)
- [Monitoring Voice ALG MGCP on page 114](#)
- [Monitoring Voice ALG SCCP on page 117](#)
- [Monitoring Voice ALG SIP on page 120](#)
- [Monitoring Voice ALG Summary on page 125](#)

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 18 on page 101](#) summarizes key output fields in the H.323 counters display.

Table 18: 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 18: 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 3](#)
 - [Monitoring Interfaces on page 142](#)

Monitoring MGCP ALGs

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

This section contains the following topics:

- [Monitoring MGCP ALG Calls on page 103](#)
- [Monitoring MGCP ALG Counters on page 103](#)
- [Monitoring MGCP ALG Endpoints on page 105](#)

Monitoring MGCP ALG Calls

Supported Platforms SRX Series, vSRX

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 19 on page 103](#) summarizes key output fields in the MGCP calls display.

Table 19: 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

Supported Platforms SRX Series, vSRX

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 20 on page 104](#) summarizes key output fields in the MGCP counters display.

Table 20: 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

Supported Platforms [SRX Series, vSRX](#)

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 21 on page 105](#) summarizes key output fields in the MGCP endpoints display.

Table 21: 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 3](#)
- [Monitoring Interfaces on page 142](#)

Monitoring SCCP ALGs

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, vSRX](#)

This section contains the following topics:

- [Monitoring SCCP ALG Calls on page 106](#)
- [Monitoring SCCP ALG Counters on page 106](#)

Monitoring SCCP ALG Calls

Supported Platforms [SRX Series, vSRX](#)

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 22 on page 106](#) summarizes key output fields in the SCCP calls display.

Table 22: 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

Supported Platforms [SRX Series, vSRX](#)

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 23 on page 106](#) summarizes key output fields in the SCCP counters display.

Table 23: 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.	—

Table 23: Summary of Key SCCP Counters Output Fields (*continued*)

Field	Values	Additional Information
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.	—
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.	—

Table 23: Summary of Key SCCP Counters Output Fields (*continued*)

Field	Values	Additional Information
Non specific error	Number of nonspecific errors.	—

- Related Documentation**
- [Monitoring Overview on page 3](#)
 - [Monitoring Interfaces on page 142](#)

Monitoring SIP ALGs

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring SIP ALG Calls on page 108](#)
- [Monitoring SIP ALG Counters on page 109](#)
- [Monitoring SIP ALG Rate Information on page 111](#)
- [Monitoring SIP ALG Transactions on page 112](#)

Monitoring SIP ALG Calls

Supported Platforms [SRX Series, vSRX](#)

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 24 on page 108](#) summarizes key output fields in the SIP calls display.

Table 24: 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.	—

Table 24: Summary of Key SIP Calls Output Fields (*continued*)

Field	Values	Additional Information
Remote Tag	Remote tag for the SIP ALG User Agent server.	—

Monitoring SIP ALG Counters

Supported Platforms [SRX Series](#), [vSRX](#)

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 25 on page 109](#) summarizes key output fields in the SIP counters display.

Table 25: 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.

Table 25: Summary of Key SIP Counters Output Fields (*continued*)

Field	Values	Additional Information
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.
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.	—

Table 25: Summary of Key SIP Counters Output Fields (*continued*)

Field	Values	Additional Information
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

Supported Platforms SRX Series, vSRX

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 26 on page 111 summarizes key output fields in the SIP rate display.

Table 26: 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

Supported Platforms [SRX Series, vSRX](#)

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 27 on page 112](#) summarizes key output fields in the SIP transactions display.

Table 27: 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. 	—
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 3](#)
- [Monitoring Interfaces on page 142](#)

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 28 on page 112](#) summarizes key output fields in the ALG H.323 page.

Table 28: 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.

Table 28: ALG H.323 Monitoring Page (*continued*)

Field	Value	Additional Information
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		
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.	—

Table 28: ALG H.323 Monitoring Page (*continued*)

Field	Value	Additional Information
H.323 Message Counter		
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 125](#)
 - [Monitoring Voice ALG MGCP on page 114](#)
 - [Monitoring Voice ALG SCCP on page 117](#)
 - [Monitoring Voice ALG SIP on page 120](#)

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 29 on page 115 summarizes key output fields in the voice ALG MGCP page.

Table 29: 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 29: 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 29: 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 125](#)
 - [Monitoring Voice ALG H.323 on page 112](#)
 - [Monitoring Voice ALG SCCP on page 117](#)
 - [Monitoring Voice ALG SIP on page 120](#)

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 30 on page 117](#) summarizes key output fields in the voice ALG SCCP page.

Table 30: 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.
Refresh	Displays the option to refresh the page.	—

Table 30: Voice ALG SCCP Monitoring Page (*continued*)

Field	Value	Additional Information
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 30: Voice ALG SCCP Monitoring Page (*continued*)

Field	Value	Additional Information
Category	<p>Displays the following categories:</p> <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 125](#)

- [Monitoring Voice ALG H.323 on page 112](#)
- [Monitoring Voice ALG MGCP on page 114](#)
- [Monitoring Voice ALG SIP on page 120](#)

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 31 on page 120](#) summarizes key output fields in the voice ALG SIP page.

Table 31: 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 31: 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 (<i>continued</i>)		

Table 31: 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 31: 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 31: 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 125](#)
 - [Monitoring Voice ALG H.323 on page 112](#)
 - [Monitoring Voice ALG MGCP on page 114](#)
 - [Monitoring Voice ALG SCCP on page 117](#)

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 32 on page 125](#) summarizes key output fields in the voice ALG summary page.

Table 32: 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 112](#)
 - [Monitoring Voice ALG MGCP on page 114](#)
 - [Monitoring Voice ALG SCCP on page 117](#)
 - [Monitoring Voice ALG SIP on page 120](#)

CHAPTER 9

Monitoring Class of Service

- [Monitoring Class-of-Service Performance on page 127](#)
- [Monitoring CoS Classifiers on page 135](#)

Monitoring Class-of-Service Performance

Supported Platforms [SRX Series, vSRX](#)

The J-Web user interface provides information about the class-of-service (CoS) performance on a device. You can view information about the current status of CoS components—classifiers, CoS value aliases, red drop profiles, forwarding classes, rewrite rules and scheduler maps. You can also see the interfaces to which these components are assigned.

In addition, you can display the entire CoS configuration, including system-chosen defaults, by entering the **show class-of-service** command.

This section contains the following topics:

- [Monitoring CoS Interfaces on page 127](#)
- [Monitoring CoS Classifiers on page 128](#)
- [Monitoring CoS Value Aliases on page 129](#)
- [Monitoring CoS RED Drop Profiles on page 130](#)
- [Monitoring CoS Forwarding Classes on page 131](#)
- [Monitoring CoS Rewrite Rules on page 132](#)
- [Monitoring CoS Scheduler Maps on page 133](#)

Monitoring CoS Interfaces

Supported Platforms [SRX Series, vSRX](#)

Purpose Display details about the physical and logical interfaces and the CoS components assigned to them.

Action Select **Monitor>Class of Service>Interfaces** in the J-Web user interface, or enter the **show class-of-service interface *interface*** command.

[Table 33 on page 128](#) summarizes key output fields for CoS interfaces.

Table 33: Summary of Key CoS Interfaces Output Fields

Field	Values	Additional Information
Interface	Name of a physical interface to which CoS components are assigned.	To display names of logical interfaces configured on this physical interface, click the plus sign (+).
Scheduler Map	Name of the scheduler map associated with this interface.	–
Queues Supported	Number of queues you can configure on the interface.	–
Queues in Use	Number of queues currently configured.	–
Logical Interface	Name of a logical interface on the physical interface, to which CoS components are assigned.	–
Object	Category of an object—for example, classifier , scheduler-map , or rewrite .	–
Name	Name that you have given to an object—for example, ba-classifier .	–
Type	Type of an object—for example, dscp , or exp for a classifier.	–
Index	Index of this interface or the internal index of a specific object.	–

Monitoring CoS Classifiers

Supported Platforms [SRX Series](#), [vSRX](#)

Purpose Display the mapping of incoming CoS value to forwarding class and loss priority.

Action For each classifier, select **Monitor>Class of Service>Classifiers** in the J-Web user interface, or enter the **show class-of-service classifier** command.

[Table 34 on page 128](#) summarizes key output fields for CoS classifiers.

Table 34: Summary of Key CoS Classifier Output Fields

Classifier Name	Name of a classifier.	To display classifier assignments, click the plus sign (+).
-----------------	-----------------------	---

Table 34: Summary of Key CoS Classifier Output Fields (*continued*)

CoS Value Type	<p>The classifiers are displayed by type:</p> <ul style="list-style-type: none"> • dscp—All classifiers of the DSCP type. • dscp ipv6—All classifiers of the DSCP IPv6 type. • exp—All classifiers of the MPLS EXP type. • ieee-802.1—All classifiers of the IEEE 802.1 type. • inet-precedence—All classifiers of the IP precedence type.
Index	Internal index of the classifier.
Incoming CoS Value	CoS value of the incoming packets, in bits. These values are used for classification.
Assign to Forwarding Class	Forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the device.
Assign to Loss Priority	Loss priority value that the classifier assigns to the incoming packet based on its CoS value.

Monitoring CoS Value Aliases

Supported Platforms [SRX Series, vSRX](#)

Purpose Display information about the CoS value aliases that the system is currently using to represent DSCP, DSCP IPv6, MPLS EXP, and IPv4 precedence bits.

Action Select **Monitor > Class of Service > CoS Value Aliases** in the J-Web user interface, or enter the **show class-of-service code-point-aliases** command.

[Table 35 on page 130](#) summarizes key output fields for CoS value aliases.

Table 35: Summary of Key CoS Value Alias Output Fields

Field	Values	Additional Information
CoS Value Type	Type of the CoS value: <ul style="list-style-type: none"> • dscp—Examines Layer 3 packet headers for IP packet classification. • dscp ipv6—Examines Layer 3 packet headers for IPv6 packet classification. • exp—Examines Layer 2 packet headers for MPLS packet classification. • ieee-802.1—Examines Layer 2 packet header for packet classification. • inet-precedence—Examines Layer 3 packet headers for IP packet classification. 	To display aliases and bit patterns, click the plus sign (+).
CoS Value Alias	Name given to a set of bits—for example, af11 is a name for 001010 bits.	—
Bit Pattern	Set of bits associated with an alias.	—

Monitoring CoS RED Drop Profiles

Supported Platforms [SRX Series](#), [vSRX](#)

Purpose Display data point information for each CoS random early detection (RED) drop profile currently on a system.

Action Select **Monitor>Class of Service>RED Drop Profiles** in the J-Web user interface, or enter the **show class-of-service drop-profile** command.

[Table 36 on page 130](#) summarizes key output fields for CoS RED drop profiles.

Table 36: Summary of Key CoS RED Drop Profile Output Fields

Field	Values	Additional Information
RED Drop Profile Name	Name of the RED drop profile. A drop profile consists of pairs of values between 0 and 100, one for queue buffer fill level and one for drop probability, that determine the relationship between a buffer's fullness and the likelihood it will drop packets.	To display profile values, click the plus sign (+).
Graph RED Profile	Link to a graph of a RED curve that the system uses to determine the drop probability based on queue buffer fullness.	The x axis represents the queue buffer fill level, and the y axis represents the drop probability.

Table 36: Summary of Key CoS RED Drop Profile Output Fields (*continued*)

Field	Values	Additional Information
Type	Type of a specific drop profile: <ul style="list-style-type: none"> • interpolated—The two coordinates (x and y) of the graph are interpolated to produce a smooth profile. • segmented—The two coordinates (x and y) of the graph are represented by line fragments to produce a segmented profile. 	—
Index	Internal index of this drop profile.	—
Fill Level	Percentage fullness of a buffer queue. This value is the x coordinate of the RED drop profile graph.	—
Drop Probability	Drop probability of a packet corresponding to a specific queue buffer fill level. This value is the y coordinate of the RED drop profile graph.	—

Monitoring CoS Forwarding Classes

Supported Platforms [SRX Series, vSRX](#)

Purpose View the current assignment of CoS forwarding classes to queue numbers on the system.

Action Select **Monitor>Class of Service>Forwarding Classes** in the J-Web user interface, or enter the **show class-of-service forwarding-class** command.

[Table 37 on page 132](#) summarizes key output fields for CoS forwarding classes.

Table 37: Summary of Key CoS Forwarding Class Output Fields

Field	Values	Additional Information
Forwarding Class	<p>Names of forwarding classes assigned to queue numbers. By default, the following forwarding classes are assigned to queues 0 through 3:</p> <ul style="list-style-type: none"> • best-effort—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value, and RED drop profiles are more aggressive. • expedited-forwarding—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. • assured-forwarding—Provides high assurance for packets within specified service profile. Excess packets are dropped. • network-control—Packets can be delayed but not dropped. 	—
Queue	Queue number corresponding to the forwarding class name.	By default, four queues, 0 through 3, are assigned to forwarding classes.

Monitoring CoS Rewrite Rules

Supported Platforms [SRX Series, vSRX](#)

Purpose Display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

Action Select **Monitor>Class of Service>Rewrite Rules** in the J-Web user interface, or enter the **show class-of-service rewrite-rules** command.

[Table 38 on page 132](#) summarizes key output fields for CoS rewrite rules.

Table 38: Summary of Key CoS Rewrite Rules Output Fields

Field	Values	Additional Information
Rewrite Rule Name	Names of rewrite rules.	—
CoS Value Type	<p>Rewrite rule type:</p> <ul style="list-style-type: none"> • dscp—For IPv4 DiffServ traffic. • dscp-ipv6—For IPv6 DiffServ traffic. • exp—For MPLS traffic. • ieee-802.1—For Layer 2 traffic. • inet-precedence—For IPv4 traffic. 	To display forwarding classes, loss priorities, and rewritten CoS values, click the plus sign (+).
Index	Internal index for this particular rewrite rule.	—

Table 38: Summary of Key CoS Rewrite Rules Output Fields (*continued*)

Field	Values	Additional Information
Forwarding Class	Forwarding class that in combination with loss priority is used to determine CoS values for rewriting.	Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.
Loss Priority	Loss priority that in combination with forwarding class is used to determine CoS values for rewriting.	—
Rewrite CoS Value To	Value that the CoS value is rewritten to.	—

Monitoring CoS Scheduler Maps

Supported Platforms [SRX Series, vSRX](#)

Purpose Display assignments of CoS forwarding classes to schedulers.

Action Select **Monitor>Class of Service>Scheduler Maps** in the J-Web user interface, or enter the **show class-of-service scheduler-map** command.

[Table 39 on page 133](#) summarizes key output fields for CoS scheduler maps.

Table 39: Summary of Key CoS Scheduler Maps Output Fields

Field	Values	Additional Information
Scheduler Map	Name of a scheduler map.	For details, click the plus sign (+).
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.	—
Scheduler Name	Name of a scheduler.	—
Forwarding Class	Forwarding classes this scheduler is assigned to.	—
Transmit Rate	Configured transmit rate of the scheduler in bits per second (bps). The rate value can be either of the following: <ul style="list-style-type: none"> A percentage—The scheduler receives the specified percentage of the total interface bandwidth. remainder—The scheduler receives the remaining bandwidth of the interface after allocation to other schedulers. 	—

Table 39: Summary of Key CoS Scheduler Maps Output Fields (*continued*)

Field	Values	Additional Information
Rate Limit	Rate limiting configuration of the queue: <ul style="list-style-type: none"> • none—No rate limiting. • exact—The queue transmits at only the configured rate. 	—
Buffer Size	Delay buffer size in the queue or the amount of transmit delay (in milliseconds). The buffer size can be either of the following: <ul style="list-style-type: none"> • A percentage—The buffer is a percentage of the total buffer allocation. • remainder—The buffer is sized according to what remains after other scheduler buffer allocations. 	—
Priority	Scheduling priority of a queue: <ul style="list-style-type: none"> • high—Packets in this queue are transmitted first. • low—Packets in this queue are transmitted last. • medium-high—Packets in this queue are transmitted after high-priority packets. • medium-low—Packets in this queue are transmitted before low-priority packets. 	—
Drop Profiles	Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.	—
Loss Priority	Packet loss priority corresponding to a drop profile: <ul style="list-style-type: none"> • low—Packet has a low loss priority. • high—Packet has a high loss priority. • medium-low—Packet has a medium-low loss priority. • medium-high—Packet has a medium-high loss priority. 	—
Protocol	Transport protocol corresponding to a drop profile.	—
Drop Profile Name	Name of the drop profile.	—

Related Documentation

- [Monitoring Overview on page 3](#)
- [Monitoring Interfaces on page 142](#)

Monitoring CoS Classifiers

Supported Platforms [SRX Series, vSRX](#)

Purpose Display the mapping of incoming CoS value to forwarding class and loss priority.

Action For each classifier, select **Monitor>Class of Service>Classifiers** in the J-Web user interface, or enter the **show class-of-service classifier** command.

[Table 34 on page 128](#) summarizes key output fields for CoS classifiers.

Table 40: Summary of Key CoS Classifier Output Fields

Classifier Name	Name of a classifier.	To display classifier assignments, click the plus sign (+).
CoS Value Type	The classifiers are displayed by type: <ul style="list-style-type: none"> • dscp—All classifiers of the DSCP type. • dscp ipv6—All classifiers of the DSCP IPv6 type. • exp—All classifiers of the MPLS EXP type. • ieee-802.1—All classifiers of the IEEE 802.1 type. • inet-precedence—All classifiers of the IP precedence type. 	
Index	Internal index of the classifier.	
Incoming CoS Value	CoS value of the incoming packets, in bits. These values are used for classification.	
Assign to Forwarding Class	Forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the device.	
Assign to Loss Priority	Loss priority value that the classifier assigns to the incoming packet based on its CoS value.	

- Related Documentation**
- [Monitoring CoS Interfaces on page 127](#)
 - [Monitoring CoS Value Aliases on page 129](#)
 - [Monitoring CoS RED Drop Profiles on page 130](#)
 - [Monitoring CoS Forwarding Classes on page 131](#)
 - [Monitoring CoS Rewrite Rules on page 132](#)
 - [Monitoring CoS Scheduler Maps on page 133](#)

CHAPTER 10

Monitoring Interfaces and Switching Functions

- [Displaying Real-Time Interface Information on page 137](#)
- [Monitoring Address Pools on page 139](#)
- [Monitoring Ethernet Switching on page 140](#)
- [Monitoring GVRP on page 141](#)
- [Monitoring Interfaces on page 142](#)
- [Monitoring MPLS Traffic Engineering Information on page 143](#)
- [Monitoring PPP on page 149](#)
- [Monitoring PPPoE on page 149](#)
- [Monitoring Spanning Tree on page 153](#)
- [Monitoring the WAN Acceleration Interface on page 154](#)

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 41 on page 138](#) and [Table 42 on page 138](#) list the keys you use to control the display using the ***interface-name*** and **traffic** options. (The keys are not case sensitive.)

Table 41: 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 42: 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)
  Output bytes:    137411893 (3344 bps)
  Input packets:   7155064 (2 pps)
  Output packets:  636071 (1 pps)
Error statistics:
  Input errors:    0
  Input drops:    0
Current delta
[2631]
[10243]
[28]
[23]
[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 142](#)

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 43 on page 139](#) summarizes key output fields in the Address Pools page.

Table 43: Address Pools Monitoring Page

Field	Values	Additional Information
Address Pool Properties		
Address Pool Name	Displays the name of the address pool.	-

Table 43: Address Pools Monitoring Page (*continued*)

Field	Values	Additional Information
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 142](#)
 - [Threats Monitoring Report on page 208](#)

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 44 on page 141](#) summarizes the Ethernet Switching output fields.

Table 44: 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 3](#)
 - [Monitoring Interfaces on page 142](#)

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 45 on page 142](#) summarizes key output fields in the GVRP page.

Table 45: 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 140](#)
 - [Monitoring Spanning Tree on page 153](#)

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 144](#)
- [Monitoring MPLS LSP Information on page 144](#)
- [Monitoring MPLS LSP Statistics on page 145](#)
- [Monitoring RSVP Session Information on page 146](#)
- [Monitoring MPLS RSVP Interfaces Information on page 148](#)

Monitoring MPLS Interfaces

Supported Platforms [SRX Series, vSRX](#)

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 46 on page 144](#) summarizes key output fields in the MPLS interface information display.

Table 46: 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.	—

Monitoring MPLS LSP Information

Supported Platforms [SRX Series, vSRX](#)

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 47 on page 144](#) summarizes key output fields in the MPLS LSP information display.

Table 47: 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.	—

Table 47: Summary of Key MPLS LSP Information Output Fields (*continued*)

Field	Values	Additional Information
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 .	—

Monitoring MPLS LSP Statistics

Supported Platforms [SRX Series](#), [vSRX](#)

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 48 on page 146 summarizes key output fields in the MPLS LSP statistics display.

Table 48: 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.	—
Total	Total number of LSPs displayed for the particular type— ingress (inbound), egress (outbound), or transit .	—

Monitoring RSVP Session Information

Supported Platforms **SRX Series, vSRX**

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 49 on page 147](#) summarizes key output fields in the RSVP session information display.

Table 49: 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.
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 .	–

Monitoring MPLS RSVP Interfaces Information

Supported Platforms [SRX Series](#), [vSRX](#)

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 50 on page 148](#) summarizes key output fields in the RSVP interfaces information display.

Table 50: 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.	—
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).	—

- Related Documentation**
- [Configuring Ping MPLS on page 249](#)
 - [MPLS Connection Checking Overview on page 247](#)
 - [Monitoring Overview on page 3](#)
 - [Monitoring Interfaces on page 142](#)

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 3](#)
 - [Monitoring Interfaces on page 142](#)

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 51 on page 150](#) summarizes key output fields in PPPoE displays.

Table 51: 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.	—
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 51: 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 51: 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 3](#)
- [Monitoring Interfaces on page 142](#)
- [Monitoring DHCP Client Bindings on page 193](#)

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 52 on page 153 summarizes key output fields in the spanning tree page.

Table 52: 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 52: 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 140](#)
 - [Monitoring GVRP on page 141](#)

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 3](#)
 - [Monitoring Interfaces on page 142](#)

CHAPTER 11

Monitoring NAT

- [Monitoring NAT on page 155](#)

Monitoring NAT

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring Source NAT Information on page 155](#)
- [Monitoring Destination NAT Information on page 161](#)
- [Monitoring Static NAT Information on page 163](#)
- [Monitoring Incoming Table Information on page 164](#)
- [Monitoring Interface NAT Port Information on page 165](#)

Monitoring Source NAT Information

Supported Platforms [SRX Series, vSRX](#)

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 53 on page 155](#) describes the available options for monitoring source NAT.

Table 53: Source NAT Monitoring Page

Field	Description	Action
Rules		

Table 53: Source NAT Monitoring Page (*continued*)

Field	Description	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.	—
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.	—

Table 53: Source NAT Monitoring Page (*continued*)

Field	Description	Action
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 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.	—

Table 53: Source NAT Monitoring Page (*continued*)

Field	Description	Action
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		
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.	–

Table 53: Source NAT Monitoring Page (*continued*)

Field	Description	Action
Reflective protocol	Translated protocol.	—
Source NAT pool	Name of the source NAT pool where persistent NAT is used.	—
Type	Persistent NAT type.	—
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.	—

Table 53: Source NAT Monitoring Page (*continued*)

Field	Description	Action
External address	Displays the external IP address.	—
Resource Usage		
Utilization for all source pools		
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	Number of used resources in the pool. For Non-PAT pools, the number of used IP addresses is displayed. For PAT pools, the number of used ports is displayed.	—
Available	Number of available resources in the pool. For Non-PAT pools, the number of available IP addresses is displayed. For PAT pools, the number of available ports is displayed.	—
Total	Number of used and available resources in the pool. For Non-PAT pools, the total number of used and available IP addresses is displayed. For PAT pools, the total number of used and available ports is displayed.	—
Usage	Percent of resources used. For Non-PAT pools, the percent of IP addresses used is displayed. For PAT pools, the percent of ports, including single and twin ports, is displayed.	—
Peak usage	Percent of resources used during the peak date and time.	—

Table 53: Source NAT Monitoring Page (*continued*)

Field	Description	Action
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.	—
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.	—

Monitoring Destination NAT Information

Supported Platforms [SRX Series, vSRX](#)

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 54 on page 161](#) summarizes key output fields in the destination NAT display.

Table 54: 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 54: 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 54: 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.	—

Monitoring Static NAT Information

Supported Platforms [SRX Series, vSRX](#)

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 55 on page 163](#) summarizes key output fields in the static NAT display.

Table 55: 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.	—
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	—

Table 55: Summary of Key Static NAT Output Fields (*continued*)

Field	Values	Action
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.	–

Monitoring Incoming Table Information

Supported Platforms [SRX Series](#), [vSRX](#)

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 56 on page 165 summarizes key output fields in the incoming table display.

Table 56: 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.

Monitoring Interface NAT Port Information

Supported Platforms [SRX Series, vSRX](#)

Purpose View port usage for an interface source pool information.

Action Select **Monitor>Firewall/NAT>Interface NAT** in the J-Web user interface, or enter the following CLI command:

- **show security nat interface-nat-ports**

Table 57 on page 165 summarizes key output fields in the interface NAT display.

Table 57: 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.	—

Table 57: Summary of Key Interface NAT Output Fields (*continued*)

Field	Values	Additional Information
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.	—

- Related Documentation**
- [Monitoring Overview on page 3](#)
 - [Monitoring Interfaces on page 142](#)

CHAPTER 12

Monitoring Security Policies

- [Monitoring Policy Statistics on page 167](#)
- [Monitoring Routing Information on page 168](#)
- [Monitoring Security Events by Policy on page 175](#)
- [Monitoring Security Features on page 177](#)

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 302](#)
 - [Checking a Security Policy Commit Failure on page 302](#)
 - [Verifying a Security Policy Commit on page 303](#)
 - [Debugging Policy Lookup on page 303](#)

Monitoring Routing Information

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring Route Information on page 168](#)
- [Monitoring RIP Routing Information on page 170](#)
- [Monitoring OSPF Routing Information on page 171](#)
- [Monitoring BGP Routing Information on page 173](#)

Monitoring Route Information

Supported Platforms [SRX Series, vSRX](#)

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 58 on page 169 describes the different filters, their functions, and the associated actions.

Table 59 on page 169 summarizes key output fields in the routing information display.

Table 58: 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 59: 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 59: 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. 	—

Monitoring RIP Routing Information

Supported Platforms [SRX Series](#), [vSRX](#)

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 60 on page 170](#) summarizes key output fields in the RIP routing display in the J-Web user interface.

Table 60: Summary of Key RIP Routing Output Fields

Field	Values	Additional Information
RIP Statistics		
Protocol Name	The RIP protocol name.	—

Table 60: Summary of Key RIP Routing Output Fields (*continued*)

Field	Values	Additional Information
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.	–

Monitoring OSPF Routing Information

Supported Platforms [SRX Series](#), [vSRX](#)

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 61 on page 172](#) summarizes key output fields in the OSPF routing display in the J-Web user interface.

Table 61: 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.	—
Total Retransmits	Number of retransmission entries enqueued.	—
Total Database Summaries	Total number of database description packets.	—
OSPF Neighbors		
Address	Address of the neighbor.	—

Table 61: Summary of Key OSPF Routing Output Fields (*continued*)

Field	Values	Additional Information
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.	–

Monitoring BGP Routing Information

Supported Platforms [SRX Series](#), [vSRX](#)

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 62 on page 174](#) summarizes key output fields in the BGP routing display in the J-Web user interface.

Table 62: 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.	–
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.	–

Table 62: Summary of Key BGP Routing Output Fields (*continued*)

Field	Values	Additional Information
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.	—

- Related Documentation**
- [Monitoring Overview on page 3](#)
 - [Monitoring Interfaces on page 142](#)

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 1. Select **Monitor>Events and Alarms>Security Events** in the J-Web user interface. The View Policy Log pane appears. [Table 63 on page 175](#) describes the content of this pane.

Table 63: 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.

Table 63: View Policy Log Fields (*continued*)

Field	Value
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 64 on page 177](#) describes the contents of this pane.

Table 64: 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 3](#)
 - [Monitoring Interfaces on page 142](#)
 - [Monitoring Alarms on page 47](#)
 - [Monitoring Events on page 193](#)

Monitoring Security Features

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring Policies on page 177](#)
- [Checking Policies on page 180](#)
- [Monitoring Screen Counters on page 183](#)
- [Monitoring IDP Status on page 185](#)
- [Monitoring Flow Gate Information on page 186](#)
- [Monitoring Firewall Authentication Table on page 187](#)
- [Monitoring Firewall Authentication History on page 189](#)
- [Monitoring 802.1x on page 191](#)

Monitoring Policies

Supported Platforms [SRX Series, vSRX](#)

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 65 on page 178](#) 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 65: 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.
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.	—

Table 65: Security Policies Monitoring Output Fields (*continued*)

Field	Value	Additional Information
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	<p>Displays the dynamic application signatures to be used as match criteria if an application firewall rule set is configured for the policy.</p> <p>For a network firewall, a dynamic application is not defined.</p>	<p>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.</p> <p>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.</p>
Action	<p>Displays the action portion of the rule set if an application firewall rule set is configured for the policy.</p> <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.
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.

Table 65: Security Policies Monitoring Output Fields (*continued*)

Field	Value	Additional Information
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.

Checking Policies

Supported Platforms [SRX Series](#), [vSRX](#)

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 66 on page 181](#) 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 66: 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.
Source Identity	Name of the source identity.

Table 66: Check Policies Output (*continued*)

Field	Function
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.
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.

Table 66: Check Policies Output (*continued*)

Field	Function
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.

Monitoring Screen Counters

Supported Platforms [SRX Series, vSRX](#)

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 67 on page 183](#) summarizes key output fields in the screen counters display.

Table 67: 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.
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.

Table 67: Summary of Key Screen Counters Output Fields (*continued*)

Field	Values	Additional Information
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.
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.	—

Table 67: Summary of Key Screen Counters Output Fields (*continued*)

Field	Values	Additional Information
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.	—

Monitoring IDP Status

Supported Platforms [SRX Series, vSRX](#)

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, 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**

Table 68 on page 186 summarizes key output fields in the IDP display.

Table 68: 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.	—
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.	—

Monitoring Flow Gate Information

Supported Platforms [SRX Series](#), [vSRX](#)

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 69 on page 187 summarizes key output fields in the flow gate display.

Table 69: 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.	—
Reference count	Number of resource manager references to the pinhole.	—
Resource	Resource manager information about the pinhole.	—

Monitoring Firewall Authentication Table

Supported Platforms [SRX Series, vSRX](#)

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 70 on page 188 summarizes key output fields in firewall authentication table display.

Table 70: 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.	—
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.	—

Table 70: Summary of Key Firewall Authentication Table Output Fields (*continued*)

Field	Values	Additional Information
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

Supported Platforms [SRX Series](#), [vSRX](#)

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 71 on page 189](#) summarizes key output fields in firewall authentication history display.

Table 71: Summary of Key Firewall Authentication History Output Fields

Field	Values	Additional Information
History of Firewall Authentication Data		
Total authentications	Number of authentication.	—
History Table		

Table 71: Summary of Key Firewall Authentication History Output Fields (*continued*)

Field	Values	Additional Information
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.	—
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: <i>Source Ip</i>		
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.	—

Table 71: Summary of Key Firewall Authentication History Output Fields (*continued*)

Field	Values	Additional Information
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

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

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 72 on page 191 summarizes the Dot1X output fields.

Table 72: 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.	–

Table 72: Summary of Dot1X Output Fields *(continued)*

Field	Values	Additional Information
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.	–

- Related Documentation
- [Monitoring Overview on page 3](#)
 - [Monitoring Interfaces on page 142](#)

CHAPTER 13

Monitoring Events, Services and System

- [Monitoring DHCP Client Bindings on page 193](#)
- [Monitoring Events on page 193](#)
- [Monitoring the System on page 196](#)

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 73 on page 193](#) summarizes the key output fields in the DHCP client binding displays.

Table 73: 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 149](#)
- [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 74 on page 194 summarizes key output fields in the events page.

Table 74: 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.	-
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.	-

Table 74: Events Monitoring Page (*continued*)

Field	Value	Additional Information
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.	-
Time	Time that the error or event occurred.	-

- Related Documentation**
- [Monitoring Alarms on page 47](#)
 - [Monitoring Security Events by Policy on page 175](#)

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 196](#)
- [Monitoring Chassis Information on page 198](#)
- [System Health Management for Branch SRX Series Devices on page 200](#)

Monitoring System Properties for SRX Series Devices

Supported Platforms [SRX Series, vSRX](#)

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

Supported Platforms [SRX Series](#), [vSRX](#)

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`

System Health Management for Branch SRX Series Devices

Supported Platforms SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX

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 3](#)
 - [Monitoring Interfaces on page 142](#)

Monitoring Unified Threat Management Features

- [Monitoring Antivirus Scan Engine Status on page 203](#)
- [Monitoring Antivirus Scan Results on page 204](#)
- [Monitoring Antivirus Session Status on page 206](#)
- [Monitoring Content Filtering Configurations on page 207](#)
- [Monitoring Reports on page 207](#)
- [Monitoring Web Filtering Configurations on page 214](#)

Monitoring Antivirus Scan Engine Status

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX](#)

Purpose 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)
```

- Related Documentation**
- [Full Antivirus Configuration Overview](#)
 - [Monitoring Antivirus Session Status on page 206](#)
 - [Monitoring Antivirus Scan Results on page 204](#)

Monitoring Antivirus Scan Results

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX345, SRX550M, vSRX](#)

Purpose 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.

Related Documentation • [Monitoring Antivirus Session Status on page 206](#)

Monitoring Antivirus Session Status

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [SRX345](#), [SRX550M](#), [vSRX](#)

Purpose 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.

Related Documentation • [Full Antivirus Configuration Overview](#)
• [Monitoring Antivirus Scan Engine Status on page 203](#)
• [Monitoring Antivirus Scan Results on page 204](#)

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 statistics** **Monitor>Security>UTM>Content Filtering** **Monitor>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.

- Related Documentation**
- [Content Filtering Overview](#)
 - [Understanding Content Filtering Protocol Support](#)
 - [Content Filtering Configuration Overview](#)
 - [Example: Attaching Content Filtering UTM Policies to Security Policies](#)

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 208](#)
- [Traffic Monitoring Report on page 212](#)

Threats Monitoring Report

Supported Platforms [SRX Series, vSRX](#)

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 75 on page 208](#) for a description of the page content.
 - **Activities** tab. See [Table 76 on page 210](#) for a description of the page content.

Table 75: Statistics Tab Output in the Threats Report

Field	Description
General Statistics Pane	
Threat Category	<p>One of the following categories of threats:</p> <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

Table 75: Statistics Tab Output in the Threats Report (*continued*)

Field	Description
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.
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.

Table 75: Statistics Tab Output in the Threats Report (*continued*)

Field	Description
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.
Hits in current hour	Number of threats per subcategory in the last hour.

Table 76: 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.

Table 76: Activities Tab Output in the Threats Report (*continued*)

Field	Function
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.
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.

Table 76: Activities Tab Output in the Threats Report (*continued*)

Field	Function
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.
Action	Action taken in response to the threat.
Last Send Time	Last time the IDP threat was sent.

Traffic Monitoring Report

Supported Platforms [SRX Series, vSRX](#)

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 77 on page 213](#) for a description of the report.

Table 77: 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.
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.

Table 77: Traffic Report Output (*continued*)

Field	Description
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.

- Related Documentation**
- [Monitoring Overview on page 3](#)
 - [Monitoring Interfaces on page 142](#)

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 Local Web Filtering](#)
- [Understanding Integrated Web Filtering](#)

Monitoring VPNs

- [Monitoring VPNs on page 217](#)

Monitoring VPNs

Supported Platforms [SRX Series, vSRX](#)

This section contains the following topics:

- [Monitoring IKE Gateway Information on page 217](#)
- [Monitoring IPsec VPN—Phase I on page 221](#)
- [Monitoring IPsec VPN—Phase II on page 222](#)
- [Monitoring IPsec VPN Information on page 223](#)

Monitoring IKE Gateway Information

Supported Platforms [SRX Series, vSRX](#)

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 78 on page 217](#) summarizes key output fields in the IKE gateway display.

Table 78: 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.

Table 78: Summary of Key IKE SA Information Output Fields (*continued*)

Field	Values	Additional Information
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.
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.	–

Table 78: 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.
Exchange Type	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. 	—
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 78: 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 78: 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

Supported Platforms SRX Series, vSRX

Purpose View IPsec VPN Phase I information.

Action Select **Monitor>IPSec VPN>Phase I** in the J-Web user interface.

Table 79 on page 221 describes the available options for monitoring IPsec VPN-Phase I.

Table 79: 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 79: 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. 	—

Monitoring IPsec VPN—Phase II

Supported Platforms SRX Series, vSRX

Purpose View IPsec VPN Phase II information.

Action Select **Monitor>IPsec VPN>Phase II** in the J-Web user interface.

Table 80 on page 222 describes the available options for monitoring IPsec VPN-Phase II.

Table 80: 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		
ID	Index number of the SA.	—
Gateway/Port	IP address of the remote gateway/port.	—

Table 80: IPsec VPN—Phase II Monitoring Page (*continued*)

Field	Values	Additional Information
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.	—

Monitoring IPsec VPN Information

Supported Platforms [SRX Series, vSRX](#)

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 81 on page 223](#) summarizes key output fields in the IPsec VPN display.

Table 81: Summary of Key IPsec VPN Information Output Fields

Field	Values	Additional Information
IPsec Security Associations		

Table 81: Summary of Key IPsec VPN Information Output Fields (*continued*)

Field	Values	Additional Information
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 81: 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 81: 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 81: 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 81: 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 3](#)
 - [Monitoring Interfaces on page 142](#)

PART 4

Troubleshooting

- [Configuring Data Path Debugging and Trace Options on page 231](#)
- [Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits on page 247](#)
- [Using Packet Capture to Analyze Network Traffic on page 265](#)
- [Troubleshooting Security Devices on page 291](#)

CHAPTER 16

Configuring Data Path Debugging and Trace Options

- [Understanding Data Path Debugging for SRX Series Devices on page 231](#)
- [Debugging the Data Path \(CLI Procedure\) on page 232](#)
- [Example: Configuring End-to-End Debugging on a High-End SRX Series Device on page 233](#)
- [Understanding Security Debugging Using Trace Options on page 237](#)
- [Setting Security Trace Options \(CLI Procedure\) on page 237](#)
- [Displaying Log and Trace Files on page 239](#)
- [Displaying Output for Security Trace Options on page 239](#)
- [Displaying Multicast Trace Operations on page 240](#)
- [Using the J-Web Traceroute Tool on page 241](#)
- [J-Web Traceroute Results and Output Summary on page 243](#)
- [Understanding Flow Debugging Using Trace Options on page 243](#)
- [Setting Flow Debugging Trace Options \(CLI Procedure\) on page 244](#)
- [Displaying a List of Devices on page 245](#)

Understanding Data Path Debugging for SRX Series Devices

Supported Platforms [SRX1500, SRX5400, SRX5600, SRX5800](#)

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 a high-end 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.

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 237](#)
- [Understanding Flow Debugging Using Trace Options on page 243](#)
- [Debugging the Data Path \(CLI Procedure\) on page 232](#)
- [Example: Configuring End-to-End Debugging on a High-End SRX Series Device on page 233](#)

Debugging the Data Path (CLI Procedure)

Supported Platforms [SRX1500](#), [SRX5400](#), [SRX5600](#), [SRX5800](#)

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
```

- 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
```

- 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 231](#)
- [Understanding Security Debugging Using Trace Options on page 237](#)
- [Understanding Flow Debugging Using Trace Options on page 243](#)
- [Setting Flow Debugging Trace Options \(CLI Procedure\) on page 244](#)

Example: Configuring End-to-End Debugging on a High-End SRX Series Device

Supported Platforms [SRX5400, SRX5600, SRX5800](#)

This example shows how to configure end-to-end debugging on an SRX Series device with an SRX5K-MPC.

- [Requirements on page 233](#)
- [Overview on page 234](#)
- [Configuration on page 234](#)
- [Enabling Data Path Debugging on page 236](#)
- [Verification on page 236](#)

Requirements

This example uses the following hardware and software components:

- SRX5600 device with an SRX5K-MPC installed with 100-Gigabit Ethernet CFP installed
- 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, then you apply an action profile. The action profile specifies a variety of actions on the processing unit. The NP ingress and NP egress are specified as location 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 200.7.6.0/24
set security datapath-debug packet-filter filter-1 destination-prefix 200.8.6.0/24
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-debug 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 200.7.6.0/24
user@host# set packet-filter filter-1 destination-prefix 200.8.6.0/24
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
```

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 200.7.6.0/24;
    destination-prefix 200.8.6.0/24;
    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 | <p>After configuring data path debugging, you must start the process on the device from operational mode.</p> <ol style="list-style-type: none">1. Enable data path debugging.

user@host> request security datapath-debug capture start

datapath-debug capture started on file datapcap2. Once you are done, you must disable data path debugging before you verify the configuration and view the reports.

user@host> request security datapath-debug capture stop

datapath-debug capture succesfully stopped, use show security datapath-debug capture to view |
|-------------------------------|---|

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.

Related Documentation

- [Understanding Data Path Debugging for SRX Series Devices on page 231](#)

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 231](#)
- [Understanding Flow Debugging Using Trace Options on page 243](#)
- [Setting Security Trace Options \(CLI Procedure\) on page 237](#)
- [Debugging the Data Path \(CLI Procedure\) on page 232](#)
- [Displaying Output for Security Trace Options on page 239](#)

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 237](#)
 - [Displaying Output for Security Trace Options on page 239](#)

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 245](#)
 - [Displaying Real-Time Monitoring Information on page 97](#)

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 237](#)
 - [Setting Security Trace Options \(CLI Procedure\) on page 237](#)

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 82 on page 240](#) summarizes the output fields of the display.

Table 82: 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.
by	IP address of the host issuing the query.
resp to address	address —Response destination address.
qid qid	qid —Query ID number.
packet from source to destination	<ul style="list-style-type: none"> • source—IP address of the source of the query or response. • destination—IP address of the destination of the query or response.

Table 82: CLI mtrace monitor Command Output Summary (*continued*)

Field	Description
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 241](#)
 - [J-Web Traceroute Results and Output Summary on page 243](#)

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 83 on page 241](#)).

Table 83: 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.

Table 83: Traceroute Field Summary (*continued*)

Field	Function	Your Action
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 4](#)
- [J-Web Traceroute Results and Output Summary on page 243](#)
- [Using the J-Web Ping MPLS Tool on page 255](#)
- [Using the J-Web Ping Host Tool on page 252](#)
- [Using the J-Web Packet Capture Tool on page 285](#)

J-Web Traceroute Results and Output Summary

Supported Platforms SRX Series, vSRX

Table 84 on page 243 summarizes the output in the traceroute display.

Table 84: 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 4](#)
- [Using the J-Web Traceroute Tool on page 241](#)

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 231](#)
- [Understanding Security Debugging Using Trace Options on page 237](#)
- [Setting Flow Debugging Trace Options \(CLI Procedure\) on page 244](#)
- [Debugging the Data Path \(CLI Procedure\) on page 232](#)

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 243](#)

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 85 on page 245 describes the **traceroute** command options.

Table 85: 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.
<i>routing-instance routing-instance-name</i>	(Optional) Uses the routing instance you specify for the traceroute.
<i>source address</i>	(Optional) Uses the source address that you specify, in the traceroute packet.
<i>tos number</i>	(Optional) Sets the type-of-service (TOS) value in the IP header of the traceroute packet. Specify a value from 0 through 255.
<i>tll number</i>	(Optional) Sets the time-to-live (TTL) value for the traceroute packet. Specify a hop count from 0 through 128.
<i>wait seconds</i>	(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 239](#)

Using MPLS to Diagnose LSPs, VPNs, and Layer 2 Circuits

- [MPLS Connection Checking Overview on page 247](#)
- [Configuring Ping MPLS on page 249](#)
- [Using the ping Command on page 250](#)
- [Using the J-Web Ping Host Tool on page 252](#)
- [J-Web Ping Host Results and Output Summary on page 254](#)
- [Using the J-Web Ping MPLS Tool on page 255](#)
- [J-Web Ping MPLS Results and Output Summary on page 258](#)
- [Pinging Layer 2 Circuits on page 259](#)
- [Pinging Layer 2 VPNs on page 260](#)
- [Pinging Layer 3 VPNs on page 261](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 262](#)

MPLS Connection Checking Overview

Supported Platforms [SRX1500, 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 86 on page 248](#) 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 86: 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 86: 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 4](#)
 - [Configuring Ping MPLS on page 249](#)
 - [Using the J-Web Ping Host Tool on page 252](#)
 - [Using the ping Command on page 250](#)

Configuring Ping MPLS

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.

- **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 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 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 device interfaces. If it is not a valid device address, the ping request fails with the error message “Can't assign requested address.”

- Related Documentation**
- [Diagnostic Tools Overview on page 4](#)
 - [MPLS Connection Checking Overview on page 247](#)
 - [Using the J-Web Ping Host Tool on page 252](#)
 - [Using the J-Web Ping MPLS Tool on page 255](#)
 - [Using the ping Command on page 250](#)

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> <ttl number> <wait seconds> <detail> <verbose>
```

[Table 87 on page 250](#) describes the **ping** command options.

To quit the **ping** command, press Ctrl-C.

Table 87: 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.
<i>countnumber</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.
<i>do-not-fragment</i>	(Optional) Sets the Don't Fragment (DF) bit in the IP header of the ping request packet.
<i>inet</i>	(Optional) Forces the ping requests to an IPv4 destination.
<i>inet6</i>	(Optional) Forces the ping requests to an IPv6 destination.
<i>interval 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.
<i>loose-source [hosts]</i>	(Optional) For IPv4, sets the loose source routing option in the IP header of the ping request packet.
<i>no-resolve</i>	(Optional) Suppresses the display of the hostnames of the hops along the path.

Table 87: CLI ping Command Options (*continued*)

Option	Description
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 .
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 4](#)
 - [Configuring Ping MPLS on page 249](#)
 - [Pinging Layer 2 Circuits on page 259](#)
 - [Pinging Layer 2 VPNs on page 260](#)
 - [Pinging Layer 3 VPNs on page 261](#)
 - [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 262](#)

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 250](#).)

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 88 on page 252](#)).

Table 88: 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.

Table 88: J-Web Ping Host Field Summary (*continued*)

Field	Function	Your Action
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.
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 4](#)
- [Configuring Ping MPLS on page 249](#)
- [J-Web Ping Host Results and Output Summary on page 254](#)
- [Using the J-Web Traceroute Tool on page 241](#)
- [Using the J-Web Ping MPLS Tool on page 255](#)
- [Using the J-Web Packet Capture Tool on page 285](#)

J-Web Ping Host Results and Output Summary

Supported Platforms [SRX Series, vSRX](#)

[Table 89 on page 254](#) summarizes the output in the ping host display.

Table 89: 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.
icmp_seq=0 icmp_seq= <i>number</i>	<i>number</i> —Sequence Number field of the ping response packet. You can use this value to match the ping response to the corresponding ping request.
ttl= <i>number</i>	<i>number</i> —Time-to-live hop-count value of the ping response packet.
<i>number</i> packets transmitted	<i>number</i> —Number of ping requests (probes) sent to host.
<i>percentage</i> packet loss	<i>percentage</i> —Number of ping responses divided by the number of ping requests, specified as a percentage.
round-trip min/avg/max/stddev = <i>min-time/avg-time/max-time/std-dev</i> ms	<ul style="list-style-type: none"> min-time—Minimum round-trip time (see time=time field in this table). avg-time—Average round-trip time. max-time—Maximum round-trip time. std-dev—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 4](#)
- [Configuring Ping MPLS on page 249](#)
- [Using the J-Web Ping Host Tool on page 252](#)

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 90 on page 255](#)).

Table 90: 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.

Table 90: J-Web Ping MPLS Field Summary (*continued*)

Field	Function	Your Action
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.
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.

Table 90: 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.
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.
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 4](#)
- [Configuring Ping MPLS on page 249](#)
- [J-Web Ping MPLS Results and Output Summary on page 258](#)
- [Using the J-Web Traceroute Tool on page 241](#)
- [Using the J-Web Ping Host Tool on page 252](#)
- [Using the J-Web Packet Capture Tool on page 285](#)

J-Web Ping MPLS Results and Output Summary

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

Table 91 on page 258 summarizes the output in the ping MPLS display.

Table 91: J-Web Ping MPLS Results and Output Summary

Field	Description
Exclamation point (!)	Echo reply was received.
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.
<i>time</i>	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 4](#)
- [Configuring Ping MPLS on page 249](#)
- [Using the J-Web Ping MPLS Tool on page 255](#)

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 92 on page 259 describes the **ping mpls l2circuit** command options.

Table 92: 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 <i>neighbor 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 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 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 250](#)
- [Configuring Ping MPLS on page 249](#)
- [Pinging Layer 2 VPNs on page 260](#)
- [Pinging Layer 3 VPNs on page 261](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 262](#)

- [Using the J-Web Ping Host Tool on page 252](#)

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 93 on page 260](#) describes the **ping mpls l2vpn** command options.

Table 93: 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 250](#)
- [Configuring Ping MPLS on page 249](#)
- [Pinging Layer 2 Circuits on page 259](#)
- [Pinging Layer 3 VPNs on page 261](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 262](#)
- [Using the J-Web Ping Host Tool on page 252](#)

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 94 on page 261](#) describes the **ping mpls l3vpn** command options.

Table 94: 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
!!!!
--- l3ping 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 250](#)
- [Configuring Ping MPLS on page 249](#)
- [Pinging Layer 2 Circuits on page 259](#)
- [Pinging Layer 2 VPNs on page 260](#)
- [Pinging RSVP-Signaled LSPs and LDP-Signaled LSPs on page 262](#)
- [Using the J-Web Ping Host Tool on page 252](#)

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 95 on page 262](#) describes the **ping mpls** command options.

Table 95: 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 250](#)
- [Configuring Ping MPLS on page 249](#)
- [Pinging Layer 2 Circuits on page 259](#)
- [Pinging Layer 2 VPNs on page 260](#)
- [Pinging Layer 3 VPNs on page 261](#)
- [Using the J-Web Ping Host Tool on page 252](#)

CHAPTER 18

Using Packet Capture to Analyze Network Traffic

- [Packet Capture Overview on page 265](#)
- [Example: Enabling Packet Capture on a Device on page 268](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 273](#)
- [Example: Configuring Packet Capture for Datapath Debugging on page 275](#)
- [Disabling Packet Capture on page 278](#)
- [Deleting Packet Capture Files on page 279](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 280](#)
- [Displaying Packet Headers on page 281](#)
- [Using the J-Web Packet Capture Tool on page 285](#)
- [J-Web Packet Capture Results and Output Summary on page 288](#)

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 266](#)
- [Firewall Filters for Packet Capture on page 267](#)
- [Packet Capture Files on page 267](#)
- [Analysis of Packet Capture Files on page 267](#)

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 268](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 273](#)
- [Using the J-Web Packet Capture Tool on page 285](#)

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 268](#)
- [Overview on page 268](#)
- [Configuration on page 269](#)
- [Verification on page 270](#)

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 270](#)
- [Verifying Captured Packets on page 270](#)

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 -xvvvvv
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 265](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 273](#)
- [Disabling Packet Capture on page 278](#)
- [Deleting Packet Capture Files on page 279](#)
- [Disabling Packet Capture on page 278](#)

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 272](#)
- [Overview on page 272](#)
- [Configuration on page 272](#)
- [Verification on page 273](#)

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	<p>Confirm that the configuration is working properly.</p> <p>Verify that packet capture is configured on the interface.</p>
Action	From configuration mode, enter the show interfaces fe-0/0/1 command.
Related Documentation	<ul style="list-style-type: none"> • Packet Capture Overview on page 265 • Changing Encapsulation on Interfaces with Packet Capture Configured on page 280 • Example: Configuring a Firewall Filter for Packet Capture on page 273 • Example: Enabling Packet Capture on a Device on page 268 • Deleting Packet Capture Files on page 279 • Disabling Packet Capture on page 278

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 273](#)
- [Overview on page 273](#)
- [Configuration on page 274](#)
- [Verification on page 275](#)

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 265](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Enabling Packet Capture on a Device on page 268](#)
- [Deleting Packet Capture Files on page 279](#)
- [Disabling Packet Capture on page 278](#)

Example: Configuring Packet Capture for Datapath Debugging

Supported Platforms [SRX1500, SRX5400, SRX5600, SRX5800](#)

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 275](#)
- [Overview on page 275](#)
- [Configuration on page 276](#)
- [Verification on page 277](#)

Requirements

Before you begin, see “[Debugging the Data Path \(CLI Procedure\)](#)” on page 232.

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.

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 277](#)
- [Verifying Data Path Debugging Capture on page 278](#)
- [Verifying Data Path Debugging Counter on page 278](#)

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 265](#)
- [Understanding Data Path Debugging for SRX Series Devices on page 231](#)
- [Debugging the Data Path \(CLI Procedure\) on page 232](#)

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 265](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 273](#)

- [Example: Enabling Packet Capture on a Device on page 268](#)
- [Deleting Packet Capture Files on page 279](#)

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 278](#)).
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 268](#)).
4. If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- [Packet Capture Overview on page 265](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 273](#)
- [Example: Enabling Packet Capture on a Device on page 268](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 280](#)
- [Disabling Packet Capture on page 278](#)

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 278](#)).
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 268](#)).
7. If you are done configuring the device, enter **commit** from configuration mode.

Related Documentation

- [Packet Capture Overview on page 265](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)
- [Example: Configuring a Firewall Filter for Packet Capture on page 273](#)

- [Example: Enabling Packet Capture on a Device on page 268](#)

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 96 on page 281](#) describes the **monitor traffic** command options.

Table 96: CLI monitor traffic Command Options

Option	Description
absolute-sequence	(Optional) Displays the absolute TCP sequence numbers.
count number	(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.
interface interface-name	(Optional) Displays packet headers for traffic on the specified interface. If an interface is not specified, the lowest numbered interface is monitored.
layer2-headers	(Optional) Displays the link-layer packet header on each line.
matching "expression"	(Optional) Displays packet headers that match an expression enclosed in quotation marks (" "). Table 97 on page 283 through Table 99 on page 284 list match conditions, logical operators, and arithmetic, binary, and relational operators you can use in the expression.
no-domain-names	(Optional) Suppresses the display of the domain name portion of the hostname.
no-promiscuous	<p>(Optional) Specifies <i>not</i> to place the monitored interface in promiscuous mode.</p> <p>In promiscuous mode, the interface reads every packet that reaches it. In nonpromiscuous mode, the interface reads only the packets addressed to it.</p>

Table 96: CLI monitor traffic Command Options (*continued*)

Option	Description
no-resolve	(Optional) Suppresses the display of hostnames.
no-timestamp	(Optional) Suppresses the display of packet header timestamps.
print-ascii	(Optional) Displays each packet header in ASCII format.
print-hex	(Optional) Displays each packet header, except link-layer headers, in hexadecimal format.
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 97 on page 283](#), enclosed in quotation marks (" "). You can combine match conditions by using the logical operators listed in [Table 98 on page 284](#) (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 99 on page 284](#) (listed from highest to lowest precedence):

- Arithmetic—Expressions that use the arithmetic operators listed in [Table 99 on page 284](#).
- Binary—Expressions that use the binary operators listed in [Table 99 on page 284](#).
- Packet data accessor—Expressions that use the following syntax:

```
protocol [byte-offset <size>]
```

Replace **protocol** with any protocol in [Table 97 on page 283](#). 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 97: CLI monitor traffic Match Conditions

Match Condition	Description
Entity Type	
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.

Table 97: CLI monitor traffic Match Conditions (*continued*)

Match Condition	Description
ip [broadcast multicast]	Matches broadcast or multicast IP packets.
ip protocol [address (\icmp igmp \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 98: 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 99: 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.

Table 99: CLI monitor traffic Arithmetic, Binary, and Relational Operators (*continued*)

Operator	Description
	Bitwise inclusive OR.
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 265](#)
- [Using the J-Web Packet Capture Tool on page 285](#)
- [Changing Encapsulation on Interfaces with Packet Capture Configured on page 280](#)
- [Example: Configuring Packet Capture on an Interface on page 271](#)

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 100 on page 286](#)). 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 100: 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 .
Addresses	Specifies the addresses to be matched for capturing the packets using a combination of the following parameters: <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>	Select address-matching criteria. For example: <ol style="list-style-type: none"> 1. From the Direction list, select source. 2. From the Type list, select host. 3. In the Address box, type 10.1.40.48. 4. Click Add.

Table 100: Packet Capture Field Summary (*continued*)

Field	Function	Your Action
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.	Select a direction and a port. For example: <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.
Header Expression	<p>Specifies the match condition for the packets to capture.</p> <p>The match conditions you specify for Addresses, Protocols, and Ports appear in expression format in this field.</p>	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 .

Table 100: Packet Capture Field Summary (*continued*)

Field	Function	Your Action
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	<p>Writes the captured packets to a file in PCAP format in <code>/var/tmp</code>. The files are named with the prefix <code>jweb-pcap</code> and the extension <code>.pcap</code>.</p> <p>If you select this option, the decoded packet headers do not appear on the packet capture page.</p>	<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 265](#)
- [Diagnostic Tools Overview on page 4](#)
- [J-Web Packet Capture Results and Output Summary on page 288](#)
- [Using the J-Web Ping MPLS Tool on page 255](#)
- [Using the J-Web Ping Host Tool on page 252](#)
- [Using the J-Web Traceroute Tool on page 241](#)

J-Web Packet Capture Results and Output Summary

Supported Platforms [SRX Series, vSRX](#)

[Table 101 on page 288](#) summarizes the output in the packet capture display.

Table 101: J-Web Packet Capture Results and Output Summary

Field	Description
<i>timestamp</i>	<p>Time when the packet was captured. The timestamp <code>00:45:40.823971</code> 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>	Direction of the packet. Specifies whether the packet originated from the Routing Engine (Out), or was destined for the Routing Engine (In).
<i>protocol</i>	<p>Protocol for the packet.</p> <p>In the sample output, IP indicates the Layer 3 protocol.</p>

Table 101: J-Web Packet Capture Results and Output Summary (*continued*)

Field	Description
source address	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. NOTE: When a string is defined for the port, the packet capture output displays the string instead of the port number.
destination address	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. NOTE: When a string is defined for the port, the packet capture output displays the string instead of the port number.
protocol	Protocol for the packet. In the sample output, TCP indicates the Layer 4 protocol.
data size	Size of the packet (in bytes).

- Related Documentation**
- [Packet Capture Overview on page 265](#)
 - [Diagnostic Tools Overview on page 4](#)
 - [Using the J-Web Packet Capture Tool on page 285](#)

Troubleshooting Security Devices

- [Recovering the Root Password for SRX Series Devices on page 291](#)
- [Troubleshooting DNS Name Resolution in Logical System Security Policies \(Master Administrators Only\) on page 292](#)
- [Troubleshooting the Link Services Interface on page 293](#)
- [Troubleshooting Security Policies on page 302](#)
- [Understanding Log Error Messages for Troubleshooting ISSU-Related Problems on page 304](#)

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`

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 [SRX1500, 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 293](#)
- [Determine What Causes Jitter and Latency on the Multilink Bundle on page 295](#)
- [Determine If LFI and Load Balancing Are Working Correctly on page 295](#)
- [Determine Why Packets Are Dropped on a PVC Between a Juniper Networks Device and a Third-Party Device on page 302](#)

Determine Which CoS Components Are Applied to the Constituent Links

Supported Platforms SRX1500, SRX300, SRX320, SRX340, vSRX

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 102 on page 294](#) shows the CoS components to be applied on a multilink bundle and its constituent links.

Table 102: 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.
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.

Table 102: CoS Components Applied on Multilink Bundles and Constituent Links (*continued*)

Cos Component	Multilink Bundle	Constituent Links	Explanation
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.

Determine What Causes Jitter and Latency on the Multilink Bundle

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [vSRX](#)

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).

Determine If LFI and Load Balancing Are Working Correctly

Supported Platforms [SRX1500](#), [SRX300](#), [SRX320](#), [SRX340](#), [vSRX](#)

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 3 on page 298 shows the overhead added to PPP and MLPPP headers.

Figure 3: 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 103 on page 298 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 103: 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 103: 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 104 on page 301](#) 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 104: 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:

- a. 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.
 - b. 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.
 - c. 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.
4. Use the results to verify load balancing.

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.

Troubleshooting Security Policies

Supported Platforms [SRX Series, vSRX](#)

- [Checking a Security Policy Commit Failure on page 302](#)
- [Verifying a Security Policy Commit on page 303](#)
- [Debugging Policy Lookup on page 303](#)

Checking a Security Policy Commit Failure

Supported Platforms [SRX Series, vSRX](#)

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.

Verifying a Security Policy Commit

Supported Platforms [SRX Series, vSRX](#)

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.

Debugging Policy Lookup

Supported Platforms [SRX Series](#)

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>`

- Related Documentation**
- [Synchronizing Policies Between Routing Engine and Packet Forwarding Engine](#)
 - [Checking a Security Policy Commit Failure on page 302](#)
 - [Verifying a Security Policy Commit on page 303](#)
 - [Debugging Policy Lookup on page 303](#)
 - [Monitoring Policy Statistics on page 167](#)

Understanding Log Error Messages for Troubleshooting ISSU-Related Problems

Supported Platforms [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 *Junos OS System Log Reference*.

- [Chassisd Process Errors on page 304](#)
- [Kernel State Synchronization on page 304](#)
- [Installation Related Errors on page 304](#)
- [ISSU Support Related Errors on page 305](#)
- [Redundancy Group Failover Errors on page 305](#)
- [Initial Validation Checks Fail on page 305](#)

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 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 ISSU.

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 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-groupss 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 following error messages are displayed when initial validation checks fail when the image is not present and 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

```
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, ISSU aborts and error messages are displayed.

Related Documentation

- [Understanding the Low-Impact ISSU Process on Devices in a Chassis Cluster](#)
- [ISSU System Requirements](#)
- [Upgrading Both Devices in a Chassis Cluster Using an ISSU](#)
- [Troubleshooting Chassis Cluster ISSU-Related Problems](#)

PART 5

Configuration Statements and Operational Commands

- [Configuration Statements on page 309](#)
- [Operational Commands on page 369](#)

CHAPTER 20

Configuration Statements

- [Accounting-Options Configuration Statement Hierarchy on page 310](#)
- [\[edit security alarms\] Hierarchy Level on page 312](#)
- [\[edit security datapath-debug\] Hierarchy Level on page 313](#)
- [\[edit security traceoptions\] Hierarchy Level on page 314](#)
- [accounting-options on page 314](#)
- [action-profile on page 315](#)
- [archive-sites on page 316](#)
- [capture-file \(Security\) on page 317](#)
- [class-usage-profile on page 318](#)
- [cluster \(Chassis\) on page 319](#)
- [counters on page 320](#)
- [datapath-debug on page 321](#)
- [decryption-failures on page 322](#)
- [destination-classes on page 323](#)
- [destination-interface on page 324](#)
- [destination-port on page 325](#)
- [fields \(for Interface Profiles\) on page 326](#)
- [fields \(for Routing Engine Profiles\) on page 327](#)
- [file \(Associating with a Profile\) on page 328](#)
- [file \(Configuring a Log File\) on page 329](#)
- [files on page 330](#)
- [filter-profile on page 331](#)
- [flow \(Security Flow\) on page 332](#)
- [global-threshold on page 334](#)
- [global-weight on page 335](#)
- [hardware-timestamp on page 335](#)
- [icmp on page 336](#)
- [idp \(Security Alarms\) on page 336](#)

- [inet6-options \(Services\) on page 337](#)
- [interface-profile on page 338](#)
- [interval on page 339](#)
- [ip-monitoring on page 340](#)
- [ip-monitoring \(Services\) on page 341](#)
- [maximum-capture-size \(Datapath Debug\) on page 342](#)
- [mib-profile on page 343](#)
- [mpls \(Security Forwarding Options\) on page 344](#)
- [next-hop on page 344](#)
- [nonpersistent on page 345](#)
- [object-names on page 345](#)
- [operation on page 346](#)
- [packet-capture on page 347](#)
- [packet-filter on page 348](#)
- [probe on page 349](#)
- [probe-interval on page 350](#)
- [probe-limit on page 351](#)
- [probe-server on page 352](#)
- [probe-type on page 353](#)
- [redundancy-group \(Chassis Cluster\) on page 354](#)
- [retry-interval \(Chassis Cluster\) on page 355](#)
- [routing-engine-profile on page 356](#)
- [rpm \(Services\) on page 357](#)
- [Security Configuration Statement Hierarchy on page 358](#)
- [size on page 360](#)
- [source-classes on page 360](#)
- [start-time on page 361](#)
- [target \(Services RPM\) on page 362](#)
- [thresholds on page 363](#)
- [traceoptions \(Security Datapath Debug\) on page 365](#)
- [transfer-interval on page 366](#)
- [traps on page 367](#)

Accounting-Options Configuration Statement Hierarchy

Supported Platforms [SRX Series, vSRX](#)

Use the statements in the **accounting-options** configuration hierarchy to collect and log data about basic system operations and services on the device.

```

accounting-options {
  class-usage-profile profile-name {
    destination-classes {
      destination-class-name;
    }
    file filename;
    interval minutes;
    source-classes {
      source-class-name;
    }
  }
  file filename {
    archive-sites url {
      password password ;
    }
    files number;
    nonpersistent;
    size bytes;
    start-time yyyy-mm-dd.hh:mm;
    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 minutes;
    object-names {
      mib-object-name;
    }
    operation (get | get-next | walk) ;
  }
  periodic-refresh disable;
  routing-engine-profile profile-name {
    fields {
      field-name ;
    }
    file filename;
    interval minutes;
  }
}

```

Related Documentation • *Administration Guide for Security Devices*

[edit security alarms] Hierarchy Level

Supported Platforms [SRX Series, vSRX](#)

```
security {
  alarms {
    audible {
      continuous;
    }
    potential-violation {
      authentication failures;
      cryptographic-self-test;
      decryption-failures {
        threshold value;
      }
      encryption-failures {
        threshold value;
      }
      idp;
      ike-phase1-failures {
        threshold value;
      }
      ike-phase2-failures {
        threshold value;
      }
      key-generation-self-test;
      non-cryptographic-self-test;
      policy {
        application {
          duration interval;
          size count;
          threshold value;
        }
        destination-ip {
          duration interval;
          size count;
          threshold value;
        }
        policy match {
          duration interval;
          size count;
          threshold value;
        }
        source-ip {
          duration interval;
          size count;
          threshold value;
        }
      }
      replay-attacks {
        threshold value;
      }
      security-log-percent-full percentage;
    }
  }
}
```


}

Related Documentation • [Security Configuration Statement Hierarchy on page 358](#)

[\[edit security datapath-debug\] Hierarchy Level](#)

Supported Platforms [SRX1500](#), [SRX5400](#), [SRX5600](#), [SRX5800](#)

```

security {
  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 files-number;
  format pacp-format;
  (no-world-readable | world-readable);
  size maximum-file-size;
}
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;
}
trace-options {
  file {
    filename;
    files files-number;
    match regular-expression;
    (no-world-readable | world-readable);
    size maximum-file-size;
  }
  no-remote-trace;
}

```

```
    }  
  }  
}
```

- Related Documentation**
- [Security Configuration Statement Hierarchy on page 358](#)
 - *Understanding Logical Systems for SRX Series Services Gateways*

[edit security traceoptions] Hierarchy Level

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX](#)

```
security {  
  traceoptions {  
    file {  
      filename;  
      files number;  
      match regular-expression;  
      (no-world-readable | world-readable);  
      size maximum-file-size;  
    }  
    flag flag;  
    no-remote-trace;  
    rate-limit messages-per-second;  
  }  
}
```

accounting-options

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

Syntax `accounting-options {...}`
`}`

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*
 - [Accounting Options Configuration on page 11](#)

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 275](#)

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 on page 17](#)

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](#)

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.

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 28](#)

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	<ul style="list-style-type: none">• ip-monitoring on page 340

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 <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	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 21

datapath-debug

Supported Platforms SRX1500, SRX5400, SRX5600, SRX5800

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.
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	<ul style="list-style-type: none">• Understanding Data Path Debugging for Logical Systems

decryption-failures

Supported Platforms	SRX1500 , SRX300 , SRX320 , SRX340 , SRX550M , vSRX
Syntax	<pre>decryption-failures { threshold <i>value</i>; }</pre>
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.
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	<ul style="list-style-type: none">• IPsec VPN Overview

destination-classes


Supported Platforms	EX Series, M Series, MX 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.
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 28

destination-interface

Supported Platforms	M Series, MX Series, SRX Series, T Series, vSRX
Syntax	destination-interface <i>interface-name</i> ;
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>], [edit services rpm probe-server (tcp udp)], [edit services rpm twamp client control-connection <i>control-client-name</i>]
Release Information	Statement introduced in Junos OS Release 7.5. Support at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
Description	<p>On M Series and T Series routers, specify a services (sp-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the sp- interface and include the unit 0 family inet statement with a /32 address.</p> <p>On M Series, MX Series, and T Series routers, specify a multiservices (ms-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the ms- interface and include the unit 0 family inet statement with a /32 address.</p> <p>The inline service interface (si- interface) is a virtual physical service interface that resides on the Packet Forwarding Engine to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot. Specify a multiservices (si-) interface that adds a timestamp to TWAMP probe messages. You must also configure the rpm twamp-client or twamp-server statement on the si- interface and include the unit 0 family inet statement with a /32 address.</p> <p>To enable RPM for the extension-provider packages on the adaptive services interface, configure the object-cache-size, policy-db-size, and package statements at the [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i> adaptive-services service-package extension-provider] hierarchy level. For the extension-provider package, <i>package-name</i> in the package <i>package-name</i> statement is jservices-rpm.</p>
Options	<i>interface-name</i> —Name of the adaptive services interface.
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"> Configuring RPM Timestamping Configuring RPM Receiver Servers hardware-timestamp on page 335

- *rpm (Interfaces)*
- *Enabling RPM for the Junos OS Extension-Provider Package*

destination-port

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	<code>destination-port port;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test <i>test-name</i>], [edit services rpm twamp client control-connection <i>control-client-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. Support at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
Description	Specify the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to which a probe is sent. This statement is used only for TCP or UDP probe types. The value for the destination-port can be only 7 when you configure the destination port along with hardware timestamping. A constraint check prevents you for configuring any other value for the destination port in this case. This constraint does not apply when you are using one-way hardware timestamping along with destination-port and either probe-type udp-ping or probe-type udp-ping-timestamp .
Options	Default: The default value for the port is 862 to which the TWAMP client establishes control connection. port —The port number can be 7 or from 49,160 through 65,535.
<div>  NOTE: The specified port numbers are recommended for RPM only. </div>	
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 BGP Neighbor Discovery Through RPM</i> • <i>Configuring RPM Probes</i>

fields (for Interface Profiles)

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<pre>fields { field-name; }</pre>
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	<p><i>field-name</i>—Name of the field:</p> <ul style="list-style-type: none">• 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	<ul style="list-style-type: none">• Configuring the Interface Profile on page 17

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 32

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 17• Configuring the Filter Profile on page 20• Configuring the MIB Profile on page 30• Configuring the Routing Engine Profile on page 32

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;
  }
  files number;
  nonpersistent;
  size bytes;
  source-classes 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.

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 14](#)

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 14

filter-profile

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	<pre>filter-profile <i>profile-name</i> { counters { counter-name; } file <i>filename</i>; interval <i>minutes</i>; }</pre>
Hierarchy Level	[edit accounting-options]
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	<p>Create a profile to filter and collect packet and byte count statistics and write them to a file in the <code>/var/log</code> directory. To apply the profile to a firewall filter, you include the accounting-profile statement at the [edit firewall filter <i>filter-name</i>] hierarchy level. For more information about firewall filters, see Firewall Filters Feature Guide for Routing Devices.</p>
Options	<p><i>profile-name</i>—Name of the filter profile.</p> <p>The remaining statements are explained separately.</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"> Configuring the Filter Profile on page 20

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;
    bpdu-vlan-flooding;
    bypass-non-ip-unicast;
    no-packet-flooding {
      no-trace-route;
    }
  }
  force-ip-reassembly;
  ipsec-performance-acceleration;
  load distribution {
    session-affinity ipsec;
  }
  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	<p>Determine how the device manages packet flow. The device can regulate packet flow in the following ways:</p> <ul style="list-style-type: none"> • Enable or disable DNS replies when there is no matching DNS request. • Set the initial session-timeout values.
Options	The remaining statements are explained separately. See CLI Explorer .
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"> • Juniper Networks Devices Processing Overview • Understanding Session Characteristics for SRX Series Services Gateways • Understanding Flow in Logical Systems for SRX Series Devices

global-threshold

Supported Platforms	SRX Series, vSRX
Syntax	global-threshold <i>number</i> ;
Hierarchy Level	[edit chassis cluster redundancy-group <i>group-number</i> 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	<i>number</i> —Value at which the IP monitoring weight will be 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	<ul style="list-style-type: none">• ip-monitoring on page 340

global-weight

Supported Platforms	SRX Series , vSRX
Syntax	<code>global-weight <i>number</i>;</code>
Hierarchy Level	[edit chassis cluster redundancy-group <i>group-number</i> 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	<p><i>number</i> —Combined weight assigned to all monitored IP addresses. A higher weight value indicates a greater importance.</p> <p>Range: 0 through 255</p> <p>Default: 255</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"> • ip-monitoring on page 340

hardware-timestamp

Supported Platforms	EX Series , MX Series , SRX Series , vSRX
Syntax	<code>hardware-timestamp;</code>
Hierarchy Level	[edit services rpm probe <i>owner</i> test <i>test-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement applied to MX Series routers in Junos OS Release 10.0.</p> <p>Statement introduced in Junos OS Release 10.3 for EX Series switches.</p>
Description	Enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with <code>icmp-ping</code> , <code>icmp-ping-timestamp</code> , <code>udp-ping</code> , and <code>udp-ping-timestamp</code> probe types.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

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

idp (Security Alarms)

Supported Platforms	SRX Series , vSRX
Syntax	<pre>idp;</pre>
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 Level	<p>security—To view this statement in the configuration.</p> <p>security-control—To add this statement to the configuration.</p>

inet6-options (Services)

Supported Platforms	vSRX
Syntax	<pre>inet6-options { source-address <i>address</i>; }</pre>
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1X49-D10 for vSRX.
Description	Specify the source IPv6 address used for probes. If the source IPv6 address is not one of the devices' assigned addresses, the packet will use the outgoing interface's address as its source.
Options	<p>inet6-options—Define the IPv6 protocol-related settings to be used for RPM probes</p> <p>source-address <i>ipv6-address</i>—Specify the base IPv6 address used to send the RPM probes from the client to the server (for example, ::ffff:a:b:c:d).</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"> • Configuring IPv6 RPM Probes on page 66

interface-profile

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

Syntax `interface-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 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.

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 17](#)

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 17 • Configuring the Filter Profile on page 20 • Configuring the MIB Profile on page 30 • Configuring the Routing Engine Profile on page 32

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 **family inet 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 334](#)
- [global-weight on page 335](#)

- [weight](#)
- [Example: Configuring Chassis Cluster Redundancy Group IP Address Monitoring on page 87](#)

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 {
          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.

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 336](#)

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	security—To view this statement in the configuration. security-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• System Log Messages

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.

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 30](#)

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](#)

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 on page 349](#)

nonpersistent

Supported Platforms	M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	nonpersistent;
Hierarchy Level	[edit accounting-options file <i>filename</i>]
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	<ul style="list-style-type: none"> • Configuring the Storage Location of the File on page 15

object-names

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	object-names { <i>mib-object-name</i> ; }
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	mib-object-name —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 30

operation

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series, vSRX
Syntax	operation <i>operation-name</i> ;
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 30

packet-capture

Supported Platforms	SRX Series, vSRX
Syntax	<pre>packet-capture { disable; file <i>filename</i> <files <i>number</i>> <size <i>bytes</i>> <world-readable no-world-readable>; maximum-capture-size <i>number</i>; }</pre>
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced in Junos OS Release 7.5.
Description	Configure packet capture on a device.
Options	<p>disable—Disable packet capture on the router.</p> <p>file <i>filename</i>—Name of the file to enable packet capture.</p> <ul style="list-style-type: none"> <i>number</i>—Maximum size of file. <i>no-world-readable</i>—Restrict file access to the owner. <i>world-readable</i>—Enable unrestricted file access. <p>maximum-capture-size—Configure the maximum size of capture for packets.</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"> Packet Capture Overview on page 265

packet-filter

Supported Platforms SRX1500, SRX5400, SRX5600, SRX5800, vSRX

Syntax 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*;
}

Hierarchy Level [edit security datapath-debug]

Release Information Command introduced in Junos OS Release 9.6 ; Support for IPv6 addresses for the **destination-prefix** and **source-prefix** options added in Junos OS Release 10.4. Support for IPv6 filter for the **interface** option added in Junos OS Release 10.4.

Description Set packet filter for taking the datapath-debug action. A maximum of four filters are supported at the same time.

- Options**
- **action-profile** (*profile-name* | default)—Identify the action profile to use. You can specify the name of the action profile to use or select default action profile.
 - **destination-port** (*port-range* | *protocol name*)—Specify a destination port to match TCP/UDP destination port.
 - **destination-prefix** *destination-prefix*—Specify a destination IPv4/IPv6 address prefix.
 - **interface** *logical-interface-name*—Specify a logical interface name.
 - **protocol** (*protocol-number* | *protocol-name*)—Match IP protocol type.
 - **source-port** (*port-range* | *protocol-name*)—Match TCP/UDP source port.
 - **source-prefix** *source-prefix*—Specify a source IP address prefix.

Required Privilege Level security—To view this statement in the configuration
security-control—To add this statement to the configuration.

probe

Supported Platforms EX Series, M Series, SRX Series, T Series, vSRX

Syntax

```

probe owner {
  test test-name {
    data-fill data;
    data-size size;
    destination-interface interface-name;
    destination-port port;
    dscp-code-point dscp-bits;
    hardware-timestamp;
    history-size size;
    inet6-options source-address ipv6-address;
    moving-average-size number;
    next-hop next-hop;
    one-way-hardware-timestamp;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    source-address address;
    target (url url | address ipv4-address | inet6-url url | inet6-address ipv6-address);
    test-interval interval;
    thresholds
    {
      egress-time microseconds;
      ingress-time microseconds;
      jitter-egress microseconds;
      jitter-ingress microseconds;
      jitter-rtt microseconds;
      rtt microseconds;
      std-dev-egress microseconds;
      std-dev-ingress microseconds;
      std-dev-rtt microseconds;
      successive-loss count;
      total-loss count;
    }
    traps [trap-names];
  }
}

```

Hierarchy Level [edit services rpm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.

Description Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.

Options *owner*—Specify an owner name up to 32 characters in length.

The remaining statements are explained separately.

Required Privilege	system—To view this statement in the configuration.
Level	interface-control—To add this statement to the configuration.

probe-interval

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	probe-interval <i>interval</i> ;
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test <i>test-name</i>], [edit services rpm twamp client control-connection <i>control-client-name</i> test-session <i>session-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. Support at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
Description	Specify the time to wait between sending packets, in seconds.
Options	<i>interval</i> —Number of seconds, from 1 through 255.
Required Privilege	interface—To view this statement in the configuration.
Level	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery Through RPM</i>• <i>Configuring RPM Probes</i>• <i>Two-Way Active Measurement Protocol Overview</i>

probe-limit

Supported Platforms	EX Series, M Series, MX Series, PTX Series, SRX Series, T Series
Syntax	probe-limit <i>limit</i> ;
Hierarchy Level	[edit services rpm]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Description	Configure the maximum number of concurrent probes allowed.
Options	<i>limit</i> —Maximum number of concurrent probes allowed. Range: 1 through 500 (PTX Series Packet Transport Routers only) 1 through 200 Default: 100
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>Limiting the Number of Concurrent RPM Probes</i>

probe-server

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

Syntax

```
probe-server {  
  tcp {  
    destination-interface interface-name;  
    port number;  
  }  
  udp {  
    destination-interface interface-name;  
    port number;  
  }  
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Description Specify the server to act as a receiver for the probes.

The remaining statements are explained separately.



NOTE: The `destination-interface` statement is not supported on PTX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring RPM Receiver Servers](#)

probe-type

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	probe-type type;
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test test-name]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.
Description	Specify the packet and protocol contents of a probe.
Options	<p>type—Specify one of the following probe type values:</p> <ul style="list-style-type: none"> • http-get—(Not available at the [edit services rpm bgp] hierarchy level.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL. • http-metadata-get—(Not available at the [edit services rpm bgp] hierarchy level.) Sends an HTTP get request for metadata to a target URL. • icmp-ping—Sends ICMP echo requests to a target address. • icmp-ping-timestamp—Sends ICMP timestamp requests to a target address. • tcp-ping—Sends TCP packets to a target. • udp-ping—Sends UDP packets to a target. • udp-ping-timestamp—Sends UDP timestamp requests to a target address.
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"> • <i>Configuring BGP Neighbor Discovery Through RPM</i>

redundancy-group (Chassis Cluster)

Supported Platforms [SRX Series, vSRX](#)

Syntax

```
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;  
}
```

Hierarchy Level [edit chassis cluster]

Release Information Statement introduced in Junos OS Release 9.0.

Description Define a redundancy group. Except for redundancy group 0, a redundancy group is a logical interface consisting of two physical Ethernet interfaces, one on each chassis. One interface is active, and the other is on standby. When the active interface fails, the standby interface becomes active. The logical interface is called a redundant Ethernet interface (**reth**).

Redundancy group 0 consists of the two Routing Engines in the chassis cluster and controls which Routing Engine is primary. You must define redundancy group 0 in the chassis cluster configuration.

Options *group-number* —Redundancy group identification number.
Range: 0 through 128



NOTE: The maximum number of redundancy groups is equal to the number of redundant Ethernet interfaces that you configure.

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 340](#)

retry-interval (Chassis Cluster)

Supported Platforms [SRX Series, vSRX](#)

Syntax `retry-interval interval;`

Hierarchy Level [edit chassis cluster redundancy-group *group-number* ip-monitoring]

Release Information Statement introduced in Junos OS Release 10.1.

Description Specify the ping packet send frequency (in seconds) for each IP address monitored by the redundancy group. (See **retry-count** for a related IP address monitoring configuration variable.)

Options *interval*—Pause time between each ping sent to each IP address monitored by the redundancy group.

Range: 1 to 30 seconds

Default: 1 second

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 340](#)

routing-engine-profile

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	<pre>routing-engine-profile <i>profile-name</i> { fields { <i>field-name</i>; } file <i>filename</i>; interval <i>minutes</i>; }</pre>
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 <code>/var/log</code> directory.
Options	<p><i>profile-name</i>—Name of the Routing Engine statistics profile.</p> <p>The remaining statements are explained separately.</p>
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 Routing Engine Profile on page 32

rpm (Services)

Supported Platforms [SRX Series, vSRX](#)

```
Syntax rpm {
    bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name <routing-instances routing-instance-name>;
        moving-average-size number-of-samples;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instances {
            routing-instance-name;
        }
        test-interval seconds;
    }
    probe owner {
        test test-name {
            data-fill data;
            data-size size;
            destination-interface interface-name;
            destination-port port;
            dscp-code-point dscp-bits;
            hardware-timestamp;
            history-size size;
            inet6-options {
                source-address address;
            }
            moving-average-size number;
            next-hop next-hop;
            one-way-hardware-timestamp;
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance instance-name;
            source-address address;
            target {
                address ipv4-address;
                url url;
                inet6-address ipv6-address;
                inet6-url url;
            }
            test-interval interval;
            thresholds {
                egress-time microseconds;
                ingress-time microseconds;
                jitter-egress microseconds;
                jitter-ingress microseconds;
                jitter-rtt microseconds;
                rtt microseconds;
                std-dev-egress microseconds;
            }
        }
    }
}
```

```

        std-dev-ingress microseconds;
        std-dev-rtt microseconds;
        successive-loss count;
        total-loss count;
    }
    traps [ trap-names];
}
}
probe-limit number;
probe-server {
    icmp {
        destination-interface interface-name;
    }
    tcp {
        destination-interface interface-name;
        port port-number;
    }
    udp {
        destination-interface interface-name;
        port port-number;
    }
}
}

```

Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Configure real-time performance monitoring (RPM) probes.
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	<ul style="list-style-type: none"> • icmp on page 336

Security Configuration Statement Hierarchy

Supported Platforms [SRX Series](#), [vSRX](#)

Use the statements in the **security** configuration hierarchy to configure actions, certificates, dynamic virtual private networks (VPNs), firewall authentication, flow, forwarding options, group VPNs, Intrusion Detection Prevention (IDP), Internet Key Exchange (IKE), Internet Protocol Security (IPsec), logging, Network Address Translation (NAT), public key infrastructure (PKI), policies, resource manager, rules, screens, secure shell known hosts, trace options, user identification, unified threat management (UTM), and zones. Statements that are exclusive to the SRX Series devices running Junos OS are described in this section.

Each of the following topics lists the statements at a sub-hierarchy of the **[edit security]** hierarchy.

- *[edit security address-book] Hierarchy Level*
- *[edit security alarms] Hierarchy Level on page 312*
- *[edit security alg] Hierarchy Level*
- *[edit security analysis] Hierarchy Level*
- *[edit security application-firewall] Hierarchy Level*
- *[edit security application-tracking] Hierarchy Level*
- *[edit security certificates] Hierarchy Level*
- *[edit security datapath-debug] Hierarchy Level on page 313*
- *[edit security dynamic-vpn] Hierarchy Level*
- *[edit security firewall-authentication] Hierarchy Level*
- *[edit security flow] Hierarchy Level*
- *[edit security forwarding-options] Hierarchy Level*
- *[edit security forwarding-process] Hierarchy Level*
- *[edit security gprs] Hierarchy Level*
- *[edit security group-vpn] Hierarchy Level*
- *[edit security idp] Hierarchy Level*
- *[edit security ike] Hierarchy Level*
- *[edit security ipsec] Hierarchy Level*
- *[edit security log] Hierarchy Level*
- *[edit security nat] Hierarchy Level*
- *[edit security pki] Hierarchy Level*
- *[edit security policies] Hierarchy Level*
- *[edit security resource-manager] Hierarchy Level*
- *[edit security screen] Hierarchy Level*
- *[edit security softwires] Hierarchy Level*
- *[edit security ssh-known-hosts] Hierarchy Level*
- *[edit security traceoptions] Hierarchy Level on page 314*
- *[edit security user-identification] Hierarchy Level*
- *[edit security utm] Hierarchy Level*
- *[edit security zones] Hierarchy Level*

size

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	size <i>bytes</i> ;
Hierarchy Level	[edit accounting-options <i>file 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 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: <i>x</i> to specify bytes, <i>xk</i> to specify KB, <i>xm</i> to specify MB, <i>xg</i> 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	<ul style="list-style-type: none">• Configuring the Maximum Size of the File on page 16


source-classes

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	source-classes { <i>source-class-name</i> ; }
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.
Description	Specify the source classes for which statistics are collected.
Options	source-class-name —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 28

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: YYYY-MM-DD.hh:mm
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 16

target (Services RPM)

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	target (url <i>url</i> address <i>ipv4-address</i> inet6-url <i>url</i> inet6-address <i>ipv6-address</i>);
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Packet Transport Routers. inet6-url and inet6-address options added in Junos OS Release 15.1X49-D10 for vSRX.
Description	Specify the destination address or URL used for the probes.
Options	<p>url <i>url</i>—For HTTP probe types, specify a fully formed URL that includes http:// in the URL address.</p> <p>inet6-url <i>url</i>—For HTTP probe types, specify a fully formed URL that includes http:// in the URL address. You can also specify an IPv6 address of a host in the URL to denote the destination or server to which the RPM probes must be sent.</p> <p>address <i>ipv4-address</i>—For all IPv4 probe types other than the HTTP probes, specify an IPv4 address for the target host.</p> <p>inet6-address <i>ipv6-address</i>—For all IPv6 probe types other than the HTTP probes, specify an IPv6 address for the target host.</p>
	<div> NOTE: Starting with Junos OS Release 15.1X49-D10, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address.</div>
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 IPv6 RPM Probes on page 66

thresholds

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	thresholds <i>thresholds</i> ;
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>], [edit services rpm twamp client control-connection <i>control-client-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Packet Series Transport Routers. Support at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
Description	Specify thresholds used for the probes. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded.



NOTE: If you configure a value of zero using the *thresholds* option for a certain probe parameter, the generation of SNMP traps is disabled for the corresponding probe attribute. For example, if you specify the `set thresholds jitter-egress 0` statement, it denotes that traps are not triggered when the jitter in egress time threshold is met or exceeded.

Options	<p><i>thresholds</i>—Specify one or more threshold measurements. The following options are supported:</p> <ul style="list-style-type: none"> • egress-time—Measures maximum source-to-destination time per probe. • ingress-time—Measures maximum destination-to-source time per probe. • jitter-egress—Measures maximum source-to-destination jitter per test. • jitter-ingress—Measures maximum destination-to- source jitter per test. • jitter-rtt—Measures maximum jitter per test, from 0 through 60,000,000 microseconds. • rtt—Measures maximum round-trip time per probe, in microseconds. • std-dev-egress—Measures maximum source-to-destination standard deviation per test. • std-dev-ingress—Measures maximum destination-to-source standard deviation per test. • std-dev-rtt—Measures maximum standard deviation per test, in microseconds. • successive-loss—Measures successive probe loss count, indicating probe failure. • total-loss—Measures total probe loss count indicating test failure, from 0 through 15.
---------	--

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>Configuring RPM Probes</i>• <i>Two-Way Active Measurement Protocol Overview</i>

traceoptions (Security Datapath Debug)

Supported Platforms SRX1500, SRX5400, SRX5600, SRX5800

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.

- 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 Level trace—To view this statement in the configuration.
 trace-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 *minutes*;

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 • [Configuring the Transfer Interval of the File on page 16](#)

traps

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series
Syntax	<code>traps traps;</code>
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>] [edit services rpm twamp client control-connection <i>control-client-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.3 for EX Series switches. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers. Support at the [edit services rpm twamp client control-connection <i>control-client-name</i>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
Description	Set the trap bit to generate traps for probes. Traps are sent if the configured threshold is met or exceeded.
Options	<p>traps—Specify one or more traps. The following options are supported:</p> <ul style="list-style-type: none"> • control-connection-closed—Generate traps when the control connection is closed. • egress-jitter-exceeded—Generate traps when the jitter in egress time threshold is met or exceeded. • egress-std-dev-exceeded—Generate traps when the egress time standard deviation threshold is met or exceeded. • egress-time-exceeded—Generate traps when the maximum egress time threshold is met or exceeded. • ingress-jitter-exceeded—Generate traps when the jitter in ingress time threshold is met or exceeded. • ingress-std-dev-exceeded—Generate traps when the ingress time standard deviation threshold is met or exceeded. • ingress-time-exceeded—Generate traps when the maximum ingress time threshold is met or exceeded. • jitter-exceeded—Generate traps when the jitter in round-trip time threshold is met or exceeded. • probe-failure—Generate traps when successive probe loss thresholds are crossed. • rtt-exceeded—Generate traps when the maximum round-trip time threshold is met or exceeded. • std-dev-exceeded—Generate traps when the round-trip time standard deviation threshold is met or exceeded. • test-completion—Generate traps when a test is completed. • test-failure—Generate traps when the total probe loss threshold is met or exceeded.

- **test-iteration-done**—Generate traps when all test sessions under control connections complete one test iteration.



NOTE: For RPM traps to be generated, you must configure the **remote-operations** SNMP trap category by including the **categories** statement at the **[edit snmp trap-group *trap-group-name*]** 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"> • <i>Configuring RPM Probes</i> • <i>Two-Way Active Measurement Protocol Overview</i>

CHAPTER 21

Operational Commands

- clear chassis cluster ip-monitoring failure-count
- clear chassis cluster ip-monitoring failure-count ip-address
- monitor list
- monitor start
- monitor stop
- mtrace monitor
- ping mpls l2circuit
- ping mpls l2vpn
- ping mpls l3vpn
- ping mpls ldp
- ping mpls lsp-end-point
- ping mpls rsvp
- request pppoe connect
- request pppoe disconnect
- request services ip-monitoring preempt-restore policy
- show chassis alarms
- show configuration
- show chassis cluster ip-monitoring status redundancy-group
- show interfaces (SRX Series)
- show poe interface (View)
- show poe telemetry
- show pppoe interfaces
- show pppoe 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 probe-results \(View\)](#)
- [show system alarms](#)
- [traceroute](#)

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 372](#)

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 371](#)

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
```

monitor list

Supported Platforms	EX Series, M Series, MX Series, SRX Series, T Series, vSRX
Syntax	monitor list
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 status of monitored log and trace files.
Options	This command has no options.
Additional Information	Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the syslog statement at the [edit system] hierarchy level and the options statement at the [edit routing-options] hierarchy level. The trace files generated by the routing protocol process are those configured with traceoptions statements at the [edit routing-options] , [edit interfaces] , and [edit protocols protocol] hierarchy levels.
Required Privilege Level	trace
Related Documentation	<ul style="list-style-type: none"> • monitor start on page 374 • monitor stop on page 376
List of Sample Output	monitor list on page 373
Output Fields	Table 105 on page 373 describes the output fields for the monitor list command. Output fields are listed in the approximate order in which they appear.

Table 105: monitor list Output Fields

Field Name	Field Description
monitor start	Indicates the file is being monitored.
"filename"	Name of the file that is being monitored.
Last changed	Date and time at which the file was last modified.

Sample Output

monitor list

```
user@host> monitor list
monitor start "vrrpd" (Last changed Dec 03:11:06 20)
monitor start "cli-commands" (Last changed Nov 07:3)
```

monitor start

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

Syntax `monitor start filename`

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Start displaying the system log or trace file and additional entries being added to those files.

Options *filename*—Specific log or trace file.

Additional Information Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the **syslog** statement at the **[edit system]** hierarchy level and the **options** statement at the **[edit routing-options]** hierarchy level. The trace files generated by the routing protocol process are configured with **traceoptions** statements at the **[edit routing-options]**, **[edit interfaces]**, and **[edit protocols protocol]** hierarchy levels.



NOTE: To monitor a log file within a logical system, issue the `monitor start logical-system-name/filename` command.

Required Privilege Level trace

Related Documentation

- [monitor list on page 373](#)
- [monitor stop on page 376](#)

List of Sample Output [monitor start on page 375](#)

Output Fields [Table 106 on page 374](#) describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

Table 106: monitor start Output Fields

Field Name	Field Description
<i>filename</i>	Name of the file from which entries are being displayed. This line is displayed initially and when the command switches between log files.
<i>Date and time</i>	Timestamp for the log entry.

Sample Output

monitor start

```
user@host> monitor start system-log
*** system-log***
Jul 20 15:07:34 hang sshd[5845]: log: Generating 768 bit RSA key.
Jul 20 15:07:35 hang sshd[5845]: log: RSA key generation complete.
Jul 20 15:07:35 hang sshd[5845]: log: Connection from 204.69.248.180 port 912
Jul 20 15:07:37 hang sshd[5845]: log: RSA authentication for root accepted.
Jul 20 15:07:37 hang sshd[5845]: log: ROOT LOGIN as 'root' from host.example.com
Jul 20 15:07:37 hang sshd[5845]: log: Closing connection to 204.69.248.180
```

monitor stop

Supported Platforms	EX Series , M Series , MX Series , PTX Series , T Series
Syntax	<code>monitor stop <i>filename</i></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Stop displaying the system log or trace file.
Options	<i>filename</i> —Specific log or trace file.
Additional Information	Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are those configured with the syslog statement at the [edit system] hierarchy level and the options statement at the [edit routing-options] hierarchy level. The trace files generated by the routing protocol process are those configured with traceoptions statements at the [edit routing-options] , [edit interfaces] , and [edit protocols <i>protocol</i>] hierarchy levels.
Required Privilege Level	trace
Related Documentation	<ul style="list-style-type: none">• monitor list on page 373• monitor start on page 374
List of Sample Output	monitor stop on page 376
Output Fields	This command produces no output.

Sample Output

monitor stop

```
user@host> monitor stop
```


mtrace monitor

Supported Platforms	EX Series, M Series, MX Series, QFX Series, SRX Series, T Series, vSRX
Syntax	mtrace monitor
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 11.3 for the QFX Series.
Description	Listen passively for IP multicast responses. To exit the mtrace monitor command, type Ctrl+c.
Options	none —Trace the master instance.
Required Privilege Level	view
List of Sample Output	mtrace monitor on page 378
Output Fields	Table 107 on page 377 describes the output fields for the mtrace monitor command. Output fields are listed in the approximate order in which they appear.

Table 107: mtrace monitor Output Fields

Field Name	Field Description
Mtrace query at	Date and time of the query.
by	Address of the host issuing the query.
resp to	Response destination.
qid	Query ID number.
packet from...to	IP address of the query source and default group destination.
from...to	IP address of the multicast source and the response address.
via group	IP address of the group to trace.
mxhop	Maximum hop setting.

Sample Output

mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

ping mpls l2circuit

Supported Platforms EX Series, M Series, MX Series, SRX Series, T Series, vSRX

Syntax ping mpls l2circuit (interface *interface-name* | virtual-circuit *virtual-circuit-id* neighbor *address*)
 <count *count*>
 <destination *address*>
 <detail>
 <exp *forwarding-class*>
 <logical-system (all | *logical-system-name*)>
 reply-mode (application-level-control-channel | ip-udp | no-reply)
 <size *bytes*>
 <source *source-address*>
 <sweep>
 <v1>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.

Description Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a ping mpls l2circuit command.

Options **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

interface *interface-name*—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

reply-mode—(Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel—Reply using an application level control channel.

ip-udp—Reply using an IPv4 or IPv6 UDP packet.

no-reply—Do not reply to the ping request.



NOTE: The **reply-mode** option and its suboptions **application-level-control-channel**, **ip-udp**, and **no-reply** are also available in Junos OS Release 10.2R4 and 10.3R2.

size bytes—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

vl—(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

virtual-circuit virtual-circuit-id neighbor address—Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

Additional Information You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls l2circuit interface on page 381](#)
[ping mpls l2circuit virtual-circuit detail on page 381](#)
[ping mpls l2circuit interface <interface-name> reply-mode on page 381](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms
```

ping mpls l2circuit interface <interface-name> reply-mode

```
user@host> ping mpls l2circuit interface lt-1/2/0.21 reply-mode application-level-control-channel
!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn

Supported Platforms EX Series, M Series, MX Series, SRX Series, T Series, vSRX

Syntax ping mpls l2vpn (instance *instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number* | interface *interface-name*)
 <bottom-label-ttl>
 <count *count*>
 <destination *address*>
 <detail>
 <exp forwarding-class>
 <logical-system (all | *logical-system-name*)>
 reply-mode (application-level-control-channel | ip-udp | no-reply)
 <size *bytes*>
 <source *source-address*>
 <sweep>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 The **reply-mode** option and its suboptions are introduced in Junos OS Release 10.4R1.

Description Check the operability of MPLS Layer 2 virtual private network (VPN) connections. Type Ctrl+c to interrupt a ping mpls l2vpn command.

Options **bottom-label-ttl**—(Optional) Display the time-to-live value for the bottom label in the label stack.

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp forwarding-class—(Optional) Value of the forwarding class for the MPLS ping packets.

instance *instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number*—Ping 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 ingress and egress provider edge (PE) routers or switches.

interface *interface-name*—Ping an interface configured for the Layer 2 VPN on the egress PE router or switch.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

reply-mode—(Optional) Reply mode for the ping request. This option has the following suboptions:

application-level-control-channel—Reply using an application level control channel.

ip-udp—Reply using an IPv4 or IPv6 UDP packet.

no-reply—Do not reply to the ping request.

The **reply-mode** option and its suboptions **application-level-control-channel**, **ip-udp**, and **no-reply** are also available in Junos OS Release 10.2R4 and 10.3R2.

size bytes—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls l2vpn instance on page 383](#)
[ping mpls l2vpn instance detail on page 384](#)
[ping mpls l2vpn interface <interface-name> reply-mode on page 384](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l2vpn instance

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2
```

```
!!!!!  
--- lsping statistics ---  
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn instance detail

```
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
```

ping mpls l2vpn interface <interface-name> reply-mode

```
user@host> ping mpls l2vpn interface lt-1/2/0.21 reply-mode ip-udp  
!!!!!  
--- lsping statistics ---  
5 packets transmitted, 5 packets received, 0% packet loss
```


ping mpls l3vpn

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

Syntax ping mpls l3vpn prefix *prefix-name*
 <*l3vpn-name*>
 <bottom-label-ttl>
 <count *count*>
 <destination *address*>
 <detail>
 <exp forwarding-class>
 <logical-system (all | *logical-system-name*)>
 <size *bytes*>
 <source *source-address*>
 <sweep>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a ping mpls l3vpn command.

Options **bottom-label-ttl**—(Optional) Display the time-to-live value for the bottom label in the label stack.

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp forwarding-class—(Optional) Value of the forwarding class for the MPLS ping packets.

l3vpn-name—(Optional) Layer 3 VPN name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

prefix *prefix-name*—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.

size *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (96 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter

a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

If the Layer 3 VPN traffic transits a route reflector within the network, the **ping mpls l3vpn** command does not work.

Required Privilege Level network

List of Sample Output [ping mpls l3vpn on page 386](#)
[ping mpls l3vpn detail on page 386](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
```

```
Reply for seq 5, return code: Egress-ok  
--- lsping statistics ---  
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls ldp

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

Syntax ping mpls ldp *fec*
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<instance *routing-instance-name*>
<logical-system (all | *logical-system-name*)>
<p2mp root-addr *ip-address* lsp-id *identifier*>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
size and **sweep** options introduced in Junos OS Release 9.6.
instance option introduced in Junos OS Release 10.0.
p2mp, **root-address**, and **lsp-id** options introduced in Junos OS Release 11.2.
Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Check the operability of MPLS LDP-signaled label-switched path (LSP) connections.
Type Ctrl+c to interrupt a **ping mpls** command.

Options **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

instance *routing-instance-name*—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

p2mp root-addr *ip-address* **lsp-id** *identifier*—(Optional) Ping the end points of a point-to-multipoint LSP. Enter the IP address of the point-to-multipoint LSP root and the ID number of the point-to-multipoint LSP.

size *bytes*—(Optional) Size of the LSP ping request packet (88 through 65468 bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the

router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *MPLS Applications Feature Guide for Routing Devices*.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls ldp fec count on page 389](#)
[ping mpls ldp p2mp root-addr lsp-id on page 389](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

ping mpls ldp p2mp root-addr lsp-id

```
user@host> ping mpls ldp p2mp root-addr 10.1.1.1/32 lsp-id 1 count 1
Request for seq 1, to interface 71, no label stack.
Request for seq 1, to interface 70, label 299786
Reply for seq 1, egress 10.1.1.3, return code: Egress-ok, time: 18.936 ms
Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
```

```
Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.4, return code: Egress-ok, time: 18.936 ms
Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
Reply for seq 1, egress 10.1.1.5, return code: Egress-ok, time: 18.936 ms
Local transmit time: 2009-01-12 03:50:03 PST 407.281 ms
Remote receive time: 2009-01-12 03:50:03 PST 426.217 ms
```

ping mpls lsp-end-point

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

Syntax ping mpls lsp-end-point *prefix-name*
 <count *count*>
 <destination *address*>
 <detail>
 <exp *forwarding-class*>
 <instance *routing-instance-name*>
 <logical-system (all | *logical-system-name*)>
 <size *bytes*>
 <source *source-address*>
 <sweep>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 The **instance** option was introduced in Junos OS Release 10.0.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options **count** *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination *address*—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

instance *routing-instance-name*—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

prefix-name—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.

size *bytes*—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is 88 bytes. If the endpoint is RSVP-based, the minimum size of the packet is 100 bytes. The maximum size in either case is 65468 bytes.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (lo.0).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output [ping mpls lsp-end-point detail on page 392](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

[ping mpls lsp-end-point detail](#)

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```


ping mpls rsvp

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

Syntax ping mpls rsvp
 <lsp-name>
 <count count>
 <destination address>
 <detail>
 <dynamic-bypass>
 <egress egress-address>
 <exp forwarding-class>
 <interface interface-name>
 <logical-system (all | logical-system-name)>
 <manual-bypass>
 <multipoint>
 <size bytes>
 <source source-address>
 <standby standby-path-name>
 <sweep>

Release Information Command introduced before Junos OS Release 7.4.
 The **egress** and **multipoint** options were introduced in Junos OS Release 9.2.
 The **size** and **sweep** options were introduced in Junos OS Release 9.6.
 The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2.
 Statement introduced in Junos OS Release 12.3X50 for the QFX Series.

Description Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options **count count**—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is 1 through 1,000,000. The default value is 5.

destination address—(Optional) Specify an address other than the default (127.0.0.1/32) for the ping echo requests. The address can be anything within the 127/8 subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.



NOTE: When using the **detail** option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in inaccurate one way ping trip times being reported.

In practice, it is difficult to synchronize the system times of independent Juniper Networks routers with sufficient accuracy to provide a meaningful time value for the **detail** option (even when synchronized using NTP).

dynamic-bypass—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

egress egress-address—(Optional) Only the specified egress router or switch responds to the ping request.

exp forwarding-class—(Optional) Value of the forwarding class for the MPLS ping packets.

interface—(Optional) Specify the name of the interface protected by the manual bypass LSP. This option is only available when you have also used the **manual-bypass** option.

logical-system (all | logical-system-name)—(Optional) Perform this operation on all logical systems or on the specified logical system.

lsp-name—Ping an RSVP-signaled LSP using an LSP name.

manual-bypass—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs. For this option, you must also specify the interface protected by the manual bypass LSP using the **interface** option.

multipoint—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

size bytes—(Optional) Size of the LSP ping request packet (100 through 65468 bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

source source-address—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

standby standby-path-name—(Optional) Name of the standby path.

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response might be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output

- [ping mpls rsvp \(Echo Reply Received\) on page 395](#)
- [ping mpls rsvp \(Echo Reply with Error Code\) on page 395](#)
- [ping mpls rsvp detail on page 395](#)
- [ping mpls rsvp multipoint egress detail count on page 395](#)
- [ping mpls rsvp multipoint detail count on page 395](#)
- [ping mpls rsvp destination detail count size on page 396](#)
- [ping mpls rsvp destination detail sweep size on page 396](#)

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

Sample Output

ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
loss
```

ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
Local transmit time: 1205310695s 215737us
Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 m Local transmit time:
1205310615s 347317us
```

```

Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
Local transmit time: 1205310615s 347262us
Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
Local transmit time: 1205310615s 347167us
Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

```

ping mpls rsvp destination detail count size

```

user@host> ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
    Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
    Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

```

ping mpls rsvp destination detail sweep size

```

user@router> ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
    Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
    Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
    Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
    Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
    Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
    Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
    Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
    Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms

```

```
Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

request pppoe connect

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M
Syntax	request pppoe connect
Release Information	Statement supported on SRX300, SRX320, and SRX340 is introduced in Junos OS Release 15.1X49-D60.
Description	Connect all sessions that are down.
Options	pppoe interface name— (Optional) Connect to a specified session.
Required Privilege Level	maintenance
List of Sample Output	request pppoe connect on page 398
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, SRX550M
Syntax	request pppoe disconnect
Release Information	Statement supported on SRX300, SRX320, and SRX340 is introduced in Junos OS Release 15.1X49-D60.
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
List of Sample Output	request pppoe disconnect on page 399
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 rpm probe-results \(View\) on page 468](#)
- [show services ip-monitoring status on page 464](#)

List of Sample Output [run request services ip-monitoring preempt-restore policy <policy name> on page 400](#)

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
```


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 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 473](#)

List of Sample Output [show chassis alarms on page 402](#)

Output Fields [Table 108 on page 401](#) lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

Table 108: 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.
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 configuration

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

Syntax show configuration
<statement-path>

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 configuration that currently is running on the router or switch, which is the last committed configuration.

Options none—Display the entire configuration.

statement-path—(Optional) Display one of the following hierarchies in a configuration. (Each **statement-path** option has additional suboptions not described here. See the appropriate feature guide or EX Series switch documentation for more information.)

- **access**—Network access configuration.
- **access-profile**—Access profile configuration.
- **accounting-options**—Accounting data configuration.
- **applications**—Applications defined by protocol characteristics.
- **apply-groups**—Groups from which configuration data is inherited.
- **chassis**—Chassis configuration.
- **chassis network-services**—Current running mode.
- **class-of-service**—Class-of-service configuration.
- **diameter**—Diameter base protocol layer configuration.
- **ethernet-switching-options**—(EX Series switch only) Ethernet switching configuration.
- **event-options**—Event processing configuration.
- **firewall**—Firewall configuration.
- **forwarding-options**—Options that control packet sampling.
- **groups**—Configuration groups.
- **interfaces**—Interface configuration.
- **jsrc**—JSRC partition configuration.
- **jsrc-partition**—JSRC partition configuration.
- **logical-systems**—Logical system configuration.
- **poe**—(EX Series switch only) Power over Ethernet configuration.
- **policy-options**—Routing policy option configuration.

- **protocols**—Routing protocol configuration.
- **routing-instances**—Routing instance configuration.
- **routing-options**—Protocol-independent routing option configuration.
- **security**—Security configuration.
- **services**—Service PIC applications configuration.
- **snmp**—Simple Network Management Protocol configuration.
- **system**—System parameters configuration.
- **virtual-chassis**—(EX Series switch only) Virtual Chassis configuration.
- **vlan**s—(EX Series switch only) VLAN configuration.

Additional Information The portions of the configuration that you can view depend on the user class that you belong to and the corresponding permissions. If you do not have permission to view a portion of the configuration, the text **ACCESS-DENIED** is substituted for that portion of the configuration. If you do not have permission to view authentication keys and passwords in the configuration, because the **secret** permission bit is not set for your user account, the text **SECRET-DATA** is substituted for that portion of the configuration. If an identifier in the configuration contains a space, the identifier is displayed in quotation marks.

Likewise, when you issue the **show configuration** command with the **| display set** pipe option to view the configuration as **set** commands, those portions of the configuration that you do not have permissions to view are substituted with the text **ACCESS-DENIED**.

Required Privilege Level view

Related Documentation

- *Displaying the Current Junos OS Configuration*
- *Overview of Junos OS CLI Operational Mode Commands*

List of Sample Output [show configuration on page 404](#)
[show configuration policy-options on page 405](#)

Output Fields This command displays information about the current running configuration.

Sample Output

show configuration

```
user@host> show configuration
## Last commit: 2006-10-31 14:13:00 PST by user1 version "8.2I0 [userc]"; ## last
  changed: 2006-10-31 14:05:53 PST
system {
  host-name exhost;
  domain-name ex1.net;
  backup-router 198.51.100.254;
  time-zone America/Los_Angeles;
  default-address-selection;
  name-server {
```

```

        192.0.2.254;
        192.0.2.249;
        192.0.2.176;
    }
    services {
        telnet;
    }
    tacplus-server {
        10.2.3.4 {
            secret /* SECRET-DATA */;
            ...
        }
    }
}
interfaces {
    ...
}
protocols {
    isis {
        export "direct routes";
    }
}
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}

```

show configuration policy-options

```

user@host> show configuration policy-options
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}

```

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	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> clear chassis cluster failover-count
List of Sample Output	show chassis cluster ip-monitoring status on page 407 show chassis cluster ip-monitoring status redundancy-group on page 408
Output Fields	Table 109 on page 406 lists the output fields for the show chassis cluster ip-monitoring status command.

Table 109: 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	<p>Current reachability state of the monitored IP address.</p> <p>Values for this field are: reachable, unreachable, and unknown. The status is "unknown" if Packet Forwarding Engines (PFEs) are not yet up and running.</p>
Failure count	Number of attempts to reach an IP address.
Reason	Explanation for the reported status. See Table 110 on page 407 .

Table 109: show chassis cluster ip-monitoring status Output Fields (*continued*)

Field Name	Field Description
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 110: 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 {
 <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	<ul style="list-style-type: none"> • Understanding Interfaces
List of Sample Output	show interfaces Gigabit Ethernet on page 418 show interfaces brief (Gigabit Ethernet) on page 419 show interfaces detail (Gigabit Ethernet) on page 419 show interfaces extensive (Gigabit Ethernet) on page 421 show interfaces terse on page 424 show interfaces controller (Channelized E1 IQ with Logical E1) on page 424 show interfaces controller (Channelized E1 IQ with Logical DS0) on page 424 show interfaces descriptions on page 425 show interfaces destination-class all on page 425 show interfaces diagnostics optics on page 425 show interfaces far-end-interval coc12-5/2/0 on page 426 show interfaces far-end-interval coc1-5/2/1:1 on page 426 show interfaces filters on page 427 show interfaces flow-statistics (Gigabit Ethernet) on page 427 show interfaces interval (Channelized OC12) on page 428 show interfaces interval (E3) on page 428 show interfaces interval (SONET/SDH) on page 429 show interfaces load-balancing on page 429 show interfaces load-balancing detail on page 429 show interfaces mac-database (All MAC Addresses on a Port) on page 430 show interfaces mac-database (All MAC Addresses on a Service) on page 430 show interfaces mac-database mac-address on page 431 show interfaces mc-ae on page 431 show interfaces media (SONET/SDH) on page 431 show interfaces policers on page 432 show interfaces policers interface-name on page 432 show interfaces queue on page 432 show interfaces redundancy on page 433 show interfaces redundancy (Aggregated Ethernet) on page 433 show interfaces redundancy detail on page 434 show interfaces routing brief on page 434 show interfaces routing detail on page 434 show interfaces routing-instance all on page 435 show interfaces snmp-index on page 435 show interfaces source-class all on page 435 show interfaces statistics (Fast Ethernet) on page 436 show interfaces switch-port on page 436 show interfaces transport pm on page 437 show security zones on page 438
Output Fields	<p>Table 111 on page 412 lists the output fields for the show interfaces command. Output fields are listed in the approximate order in which they appear.</p>

Table 111: 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
Current address	Configured MAC address.	detail extensive none

Table 111: show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 111: 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 statement. • 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, so for Gigabit Ethernet PICs, this number should 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 should 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 111: 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	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
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 111: 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 should 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 111: 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 111: 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: 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, Generation:
150

```

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, Runts: 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:
Total octets        Receive      Transmit
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
Jabber frames      0
Fragment frames    0
VLAN tagged frames 0
Code violations     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      Bandwidth      Buffer Priority
Limit
0 best-effort           %      bps      %      usec      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: 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,
  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      Bytes
                        (packet-per-second) (bits-per-second)
                        gold              0              0
                        (              0) (              0)
                        silver            0              0
                        (              0) (              0)
Logical interface so-0/1/3.0

      Destination class      Packets      Bytes
                        (packet-per-second) (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    inet
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 igmp 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
Is ping
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    inet
ge-0/0/0.0     up    up    inet
                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
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:
    Total bytes              Receive              Transmit
                          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
0	NA	NA	NA	NA	
Lane differential group delay	0	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	
Rx input power(0.01dBm)	-3642	-3665	-3626	-3637	-1800
-500	No	No	No	No	
Module temperature(Celsius)	46	46	46	46	-5
75	No	No	No	No	
Tx laser bias current(0.1mA)	0	0	0	0	0
0	NA	NA	NA	NA	
Rx laser bias current(0.1mA)	1270	1270	1270	1270	0
0	NA	NA	NA	NA	
Carrier frequency offset(MHz)	-186	-186	-186	-186	-5000
5000	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 poe interface (View)

Supported Platforms	SRX1500, SRX320, SRX340, SRX550M
Syntax	show poe interface <ge-fpc/pic/port>
Release Information	Command introduced in Junos OS Release 9.5.
Description	Display the status of Power over Ethernet (PoE) ports.
Options	<ul style="list-style-type: none"> none—Display the status of all PoE ports on the SRX Series device. ge-fpc/pic/port— (Optional) Display the status of a specific PoE port on the SRX Series device.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> <i>Example: Configuring PoE on All Interfaces</i>
Output Fields	Table 112 on page 440 lists the output fields for the show poe interface command. Output fields are listed in the approximate order in which they appear.

Table 112: show poe interface Output Fields

Field name	Field Description
PoE Interface	Specifies the interface name.
Admin Status	Specifies whether PoE capabilities are enabled or disabled.
Oper status	Specifies the operational status of the port.
Max-power	Specifies the maximum power configured on the port.
Priority	Specifies whether the port is high priority or low priority.
Power-consumption	Specifies how much power is being used by the port.
Class	Indicates the class of the powered device as defined by the IEEE 802 AF standard.

Sample Output

show poe interface

```
user@host>show poe interface
```

```
Interface Admin status Oper status Max power Priority Power consumption Class
ge-0/0/0 Enabled Searching 15.4W Low 0.0W 0
ge-0/0/1 Enabled Powered-up 15.4W High 6.6W 0
```


ge-0/0/2 Disabled	Disabled	15.4W	Low	0.0W	0
ge-0/0/3 Disabled	Disabled	15.4W	Low	0.0W	0

user@host>show poe interface ge-0/0/1

PoE interface status :
PoE interface : ge-0/0/1
Administrative status : Enabled
Operational status : Powered-up
Power limit on the interface : 15.4 W
Priority : High
Power consumed : 6.6 W
Class of power device : 0

show poe telemetries

Supported Platforms [SRX1500, SRX320, SRX340, SRX550M](#)

Syntax show poe telemetries
 <interface *interface-name* count *number*>
 <count *number* interface *interface-name*>

Release Information Command modified in Junos OS Release 12.3X48-D10.

Description Display a history of power consumption on the specified interface. Telemetries must be enabled on the interface before you can display a history of power consumption.

- Options**
- **Interface *interface-name***—Display telemetries for the specified PoE interface.
 - **count *number***—Display the specified number of telemetries records for the specified PoE interface.

Required Privilege Level View

Related Documentation

- *Example: Configuring PoE on All Interfaces*

Output Fields [Table 113 on page 442](#) lists the output fields for the **show poe telemetries interface** command. Output fields are listed in the approximate order in which they appear.

Table 113: show poe telemetries interface Output Fields

Field name	Field Description
S1 No	Number of the record for the specified port. The last record is the most recent.
Timestamp	Time that the power-consumption data was gathered.
Power	Amount of power provided by the specified port at the time the data was gathered.
Voltage	Voltage on the specified port at the time the data was gathered.

Sample Output

show poe telemetries interface

```
user@host>show poe telemetries interface ge-0/0/1 count 8
```

S1 No	Timestamp	Power	Voltage
1	Fri Jan 04 11:41:15 2009	6.6 W	47.2 V
2	Fri Jan 04 11:40:15 2009	6.6 W	47.2 V
3	Fri Jan 04 11:39:15 2009	6.6 W	47.2 V
4	Fri Jan 04 11:38:15 2009	6.6 W	47.2 V
5	Fri Jan 04 11:37:15 2009	6.6 W	47.2 V
6	Fri Jan 04 11:36:15 2009	6.6 W	47.2 V

```
7      Fri Jan 04 11:35:15 2009 6.6 W    47.2 V
8      Fri Jan 04 11:34:15 2009 6.6 W    47.2 V
```

user@host>show poe telemetries count 5 interface ge-0/0/1

Sl No	Timestamp	Power	Voltage
1	Fri Jan 04 11:47:15 2009	6.6 W	47.2 V
2	Fri Jan 04 11:38:15 2009	6.6 W	47.2 V
3	Fri Jan 04 11:29:15 2009	6.6 W	47.2 V
4	Fri Jan 04 11:11:15 2009	6.6 W	47.2 V
5	Fri Jan 04 11:10:15 2009	6.6 W	47.2 V

show pppoe interfaces

Supported Platforms [SRX1500, SRX300, SRX320, SRX340, SRX550M](#)

Syntax `show pppoe interfaces`
`<brief | detail | extensive>`
`<pp0.logical>`

Release Information Command introduced in Junos OS Release 9.5.

Description Display session-specific information about PPPoE interfaces.

Options **none**—Display interface information for all PPPoE interfaces.

brief | detail—(Optional) Display the specified level of output.

extensive—(Optional) Display information about the number of packets sent and received and the number of timeouts during a PPPoE session.

pp0.logical—(Optional) Name of an interface. The logical unit number for static interfaces can be a value from 0 through 16,385. The logical unit number for dynamic interfaces can be a value from 1,073,741,824 through the maximum number of logical interfaces supported on your SRX300, SRX320, and SRX340, and SRX550M devices.

Required Privilege Level view

Related Documentation

- [Understanding Ethernet Interfaces](#)

List of Sample Output [show pppoe interfaces on page 446](#)
[show pppoe interfaces brief on page 446](#)
[show pppoe interfaces detail on page 446](#)
[show pppoe interfaces extensive on page 446](#)

Output Fields [Table 114 on page 444](#) lists the output fields for the **show pppoe interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 114: show pppoe interfaces Output Fields

Field Name	Field Description
Index	Index number of the logical interface, which reflects its initialization sequence.
State	State of the logical interface: up or down .
Session ID	Session ID.
Service name	Type of service required (can be used to indicate an ISP name, a class, or quality of service).
Configured AC name	Configured access concentrator name.

Table 114: show pppoe interfaces Output Fields (*continued*)

Field Name	Field Description
Session AC name	Name of the access concentrator.
Remote MAC address or Remote MAC	MAC address of the remote side of the connection, either the access concentrator or the PPPoE client.
Auto-reconnect timeout	Timeout value for reconnecting after a PPPoE session is terminated (in seconds).
Idle timeout	Length of time (in seconds) that a connection can be idle before disconnecting.
Session uptime	Length of time the session has been up, in <i>hh:mm:ss</i> .
Ignore End-Of-List tag	Disables the End-of-List tag to continue processing of other tags after the End-of-List tag in a PPPoE Active Discovery Offer (PADO) packet.
Underlying interface	Interface on which PPPoE is running.
Packet Type	<p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <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 Termination 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 performing 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 packets—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • Unknown packets—Unrecognized packets.
Timeout	<p>Timeouts that occur during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI packets received within the timeout period. • PADO—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • PADR—No PADR packets received within the timeout period.
Receive Error Counters	<p>Error counters received during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI error counters received during the session. • PADO—No PADO error counters received during the session. • PADR—No PADR error counters received during the session. • PADS—No PADS error counters received during the session.

Sample Output

show pppoe interfaces

```
user@host> show pppoe interfaces
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

show pppoe interfaces brief

```
user@host> show pppoe interfaces brief
```

Interface	Underlying interface	State	Session ID	Remote MAC
pp0.0	ge-0/0/1.0	Session up	4	b0:c6:9a:74:5e:c1

show pppoe interfaces detail

```
user@host> show pppoe interfaces detail
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:21 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
  Ignore End-Of-List tag: Enable
```

show pppoe interfaces extensive

```
user@host> show pppoe interfaces extensive
pp0.0 Index 71
  State: Session up, Session ID: 4,
  Service name: None,
  Session AC name: srx-pppoe-ac, Configured AC name: None,
  Remote MAC address: b0:c6:9a:74:5e:c1,
  Session uptime: 5d 15:22 ago,
  Auto-reconnect timeout: Never, Idle timeout: Never,
  Underlying interface: ge-0/0/1.0 Index 70
```

PacketType	Sent	Received
PADI	1	0
PADO	0	1
PADR	1	0
PADS	0	1
PADT	0	0
Service name error	0	0
AC system error	0	0
Generic error	0	0
Malformed packets	0	0
Unknown packets	0	0
Timeout		
PADI	0	
PADO	0	
PADR	0	
Receive Error Counters		

PADI	0
PADO	0
PADR	0
PADS	0

show pppoe statistics

Supported Platforms [SRX1500, SRX300, SRX320, SRX340](#)

Syntax `show pppoe statistics`
`<logical-interface-name>`

Release Information Command is first introduced in Junos OS Release 9.5.

Description Display statistics information about PPPoE interfaces.

Options **none**—Display PPPoE statistics for all interfaces.
logical-interface-name—(Optional) Name of an underlying PPPoE logical interface.

Required Privilege Level view

Related Documentation

- [show pppoe interfaces on page 444](#)
- *Understanding Ethernet Interfaces*

List of Sample Output [show pppoe statistics on page 449](#)

Output Fields [Table 115 on page 448](#) lists the output fields for the **show pppoe statistics** command. Output fields are listed in the approximate order in which they appear.

Table 115: show pppoe statistics Output Fields

Field Name	Field Description
Active PPPoE sessions	Total number of active PPPoE sessions.
Packet Type	<p>Number of packets sent and received during the PPPoE session, categorized by packet type and packet errors:</p> <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 Termination 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 performing 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 packets—Malformed or short packets that caused the packet handler to discard the frame as unreadable. • Unknown packets—Unrecognized packets.

Table 115: show pppoe statistics Output Fields (*continued*)

Field Name	Field Description
Timeout	<p>Timeouts that occur during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI packets received within the timeout period. • PADO—No PADO packets received within the timeout period. (This value is always zero and is not supported.) • PADR—No PADR packets received within the timeout period.
Receive Error Counters	<p>Error counters received during the PPPoE session:</p> <ul style="list-style-type: none"> • PADI—No PADI error counters received during the session. • PADO—No PADO error counters received during the session. • PADR—No PADR error counters received during the session. • PADS—No PADS error counters received during the session.

Sample Output

show pppoe statistics

```

user@host> show pppoe statistics
Active PPPoE sessions: 0

PacketType          Sent      Received
PADI                0          0
PADO                0          0
PADR                0          0
PADS                0          0
PADT                0          0
Service name error  0          0
AC system error     0          0
Generic error       0          0
Malformed packets   0          0
Unknown packets     0          0
Timeout
PADI                0
PADO                0
PADR                0
Receive Error Counters
PADI                0
PADO                0
PADR                0
PADS                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 452](#)
[show security alarms detail on page 452](#)
[show security alarms alarm-id on page 452](#)
[show security alarms alarm-type authentication on page 453](#)
[show security alarms newer-than <time> on page 453](#)
[show security alarms older-than <time> on page 453](#)
[show security alarms process <process> on page 453](#)
[show security alarms severity <severity> on page 453](#)

Output Fields [Table 116 on page 451](#) 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 116: 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
Message	Information about the alarm, including the alarm type, username, IP address, and port number.	All levels

Table 116: show security alarms (*continued*)

Field Name	Field Description	Level of Output
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 [SRX1500, 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 455](#)
- [Understanding Data Path Debugging for Logical Systems](#)

List of Sample Output [show security datapath—debug capture on page 454](#)

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 [SRX1500, 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](#)
- [Understanding Data Path Debugging for Logical Systems](#)

List of Sample Output [show security datapath-debug counter on page 455](#)

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	SRX1500, SRX5400, SRX5600, SRX5800, 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.
Required Privilege Level	View
Related Documentation	<ul style="list-style-type: none"> • show security monitoring fpc fpc-number on page 458 • show security monitoring performance session on page 461 • show security monitoring performance spu on page 462

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 <i>fpc-number</i> <node (<i>node-id</i> all local primary)>
Release Information	Command introduced in Junos OS Release 9.2.
Description	Display security monitoring information about the FPC slot.
Options	<ul style="list-style-type: none"> • <i>fpc-number</i>—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. <ul style="list-style-type: none"> • <i>node-id</i>—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 Branch SRX Series Devices .
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show services ip-monitoring status on page 464
List of Sample Output	show security monitoring fpc 0 on page 459 show security monitoring fpc 1 on page 459 show security monitoring fpc 8 on page 460
Output Fields	Table 117 on page 458 lists the output fields for the show security monitoring fpc <i>fpc-number</i> command. Output fields are listed in the approximate order in which they appear.

Table 117: 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.
CPU Utilization (%)	Total percentage of CPU being used by the PIC's processors.

Table 117: show security monitoring fpc fpc-number Output Fields (*continued*)

Field Name	Field Description
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 464](#)

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 464](#)

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	<ul style="list-style-type: none"> show services rpm probe-results (View) on page 468
List of Sample Output	show services ip-monitoring status on page 465 show services ip-monitoring status on page 465 show services ip-monitoring status on page 466 show services ip-monitoring status on page 466 show services ip-monitoring status on page 466
Output Fields	Table 118 on page 464 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 118: 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.
State	Display the state of the route injection action. If the state is APPLIED, then the ip-monitoring policy is in failover state.

Table 118: show services ip-monitoring status Output Fields (*continued*)

Field Name	Field Description
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 probe-results (View)

Supported Platforms	SRX1500, SRX300, SRX320, SRX340, SRX550M, vSRX
Syntax	show services rpm probe-results <owner <i>owner</i> > <test <i>name</i> >
Release Information	Command introduced in Junos OS Release 11.2.
Description	Display the results of the most recent real-time performance monitoring (RPM) probes.
Options	<p>none—Display all results of the most recent RPM probes.</p> <p>owner <i>owner</i>—(Optional) Display information for the specified probe owner.</p> <p>test <i>name</i>—(Optional) Display information for the specified test.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show services ip-monitoring status on page 464
List of Sample Output	show services rpm probe-results (IPv4 Targets) on page 471 show services rpm probe-results (IPv6 Targets) on page 472
Output Fields	Table 119 on page 468 lists the output fields for the show services rpm probe-results command. Output fields are listed in the approximate order in which they appear.

Table 119: show services rpm probe-results Output Fields

Field Name	Field Description
Owner	Owner name. When you configure the probe owner statement at the [edit services rpm] hierarchy level, this field displays the configured owner name. When you configure BGP neighbor discovery through RPM, the output for this field is Rpm-Bgp-Owner .
Test	Name of a test representing a collection of probes. When you configure the test test-name statement at the [edit services rpm probe owner] hierarchy level, the field displays the configured test name. When you configure BGP neighbor discovery through RPM, the output for this field is Rpm-BGP-Test-<i>n</i> , where <i>n</i> is a cumulative number.
Target address	Destination IPv4 address used for the probes.
Target inet6-address	Destination IPv6 address used for the probes. This field is displayed when the probes are sent to the configured IPv6 targets or RPM servers.
Source address	Source address used for the probes.
Probe type	Protocol configured on the receiving probe server: http-get , http-metadata-get , icmp-ping , icmp6-ping , icmp-ping-timestamp , tcp-ping , udp-ping , or udp-ping-timestamp .

Table 119: show services rpm probe-results Output Fields (*continued*)

Field Name	Field Description
Test size	Number of probes within a test.
Routing Instance Name	<p>(BGP neighbor discovery) Name of the configured (if any) routing instance, logical system name, or both, in which the probe is configured:</p> <ul style="list-style-type: none"> When a routing instance is defined within a logical system, the logical system name is followed by the routing instance name. A slash (/) is used to separate the two entities. For example, if the routing instance called R1 is configured within the logical system called LS, the name in the output field is LS/R1. When a routing instance is configured but the default logical system is used, the name in the output field is the name of the routing instance. When a logical system is configured but the default routing instance is used, the name in the output field is the name of the logical system followed by default. A slash (/) is used to separate the two entities. For example, LS/default.
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.

Table 119: show services rpm probe-results Output Fields (*continued*)

Field Name	Field Description
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.
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 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 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.

Table 119: show services rpm probe-results Output Fields (*continued*)

Field Name	Field Description
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 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.

Sample Output

show services rpm probe-results (IPv4 Targets)

```
user@host> show services rpm probe-results
```

```

Owner: probe_a, Test: a1
Target address: 200.1.1.2, Probe type: icmp-ping
Destination interface name: ge-0/0/6.0
Test size: 10 probes
Probe results:
  Response received, Sat Jul 30 11:52:21 2011, No hardware timestamps
  Rtt: 1897 usec
Results over current test:
  Probes sent: 8, Probes received: 8, Loss percentage: 0
  Measurement: Round trip time
    Samples: 8, Minimum: 1897 usec, Maximum: 7205 usec, Average: 2848 usec,
    Peak to peak: 5308 usec, Stddev: 1715 usec, Sum: 22783 usec
Results over last test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0
  Test completed on Sat Jul 30 11:52:01 2011
  Measurement: Round trip time
    Samples: 10, Minimum: 1907 usec, Maximum: 8201 usec, Average: 3111 usec,

    Peak to peak: 6294 usec, Stddev: 2306 usec, Sum: 31106 usec
Results over all tests:
  Probes sent: 598, Probes received: 327, Loss percentage: 45
  Measurement: Round trip time

```

```
Samples: 327, Minimum: 1878 usec, Maximum: 133729 usec,  
Average: 3304 usec, Peak to peak: 131851 usec, Stddev: 7561 usec,  
Sum: 1080434 usec
```

show services rpm probe-results (IPv6 Targets)

```
user@host> show services rpm probe-results  
Owner: p, Test: t1  
Target inet6-address: 2001:db8:0:1:2a0:a502:0:1da,  
Target Port : 34567 Test size: 1000000 probes  
Probe results:  
  Response received, Mon May 18 10:48:07 2015, No hardware timestamps  
  Rtt: 236 usec, Round trip jitter: -10 usec, Round trip interarrival jitter:  
484 usec  
Results over current test:  
  Probes sent: 10, Probes received: 10, Loss percentage: 0  
  Measurement: Round trip time  
    Samples: 10, Minimum: 231 usec, Maximum: 298 usec, Average: 268 usec,  
Peak to peak: 67 usec, Stddev: 24 usec, Sum: 2682 usec  
  Measurement: Positive round trip jitter  
    Samples: 3, Minimum: 15 usec, Maximum: 1841 usec, Average: 750 usec, Peak  
to peak: 1826 usec, Stddev: 787 usec, Sum: 2251 usec  
  Measurement: Negative round trip jitter  
    Samples: 7, Minimum: 10 usec, Maximum: 1244 usec, Average: 709 usec, Peak  
to peak: 1234 usec, Stddev: 466 usec, Sum: 4961 usec  
Results over last test:  
  Probes sent: 10, Probes received: 10, Loss percentage: 0  
  Test completed on Mon Dec 16 10:48:07 2013  
  Measurement: Round trip time  
    Samples: 10, Minimum: 231 usec, Maximum: 298 usec, Average: 268 usec,  
Peak to peak: 67 usec, Stddev: 24 usec, Sum: 2682 usec  
  Measurement: Positive round trip jitter  
    Samples: 3, Minimum: 15 usec, Maximum: 1841 usec, Average: 750 usec, Peak  
to peak: 1826 usec, Stddev: 787 usec, Sum: 2251 usec  
  Measurement: Negative round trip jitter  
    Samples: 7, Minimum: 10 usec, Maximum: 1244 usec, Average: 709 usec, Peak  
to peak: 1234 usec, Stddev: 466 usec, Sum: 4961 usec  
Results over all tests(From start of current control session):  
  Probes sent: 490, Probes received: 488, Loss percentage: 0  
  Measurement: Round trip time  
    Samples: 488, Minimum: 231 usec, Maximum: 306 usec, Average: 270 usec,  
Peak to peak: 75 usec, Stddev: 16 usec, Sum: 131586 usec  
  Measurement: Positive round trip jitter  
    Samples: 254, Minimum: 0 usec, Maximum: 10151 usec, Average: 157 usec,  
Peak to peak: 10151 usec, Stddev: 873 usec, Sum: 39817 usec  
  Measurement: Negative round trip jitter  
    Samples: 233, Minimum: 1 usec, Maximum: 10170 usec, Average: 171 usec,  
Peak to peak: 10169 usec, Stddev: 888 usec, Sum: 39889 usec
```


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. 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 473

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

```

traceroute

Supported Platforms	EX Series, M Series, MX Series, OCX1100, PTX Series, QFabric System, QFX Series standalone switches, T Series
List of Syntax	Syntax on page 474 Syntax (QFX Series and OCX Series) on page 474
Syntax	<pre>traceroute <i>host</i> <as-number-lookup> <bypass-routing> <clns> <gateway <i>address</i>> <inet inet6> <interface <i>interface-name</i>> <logical system <i>logical-system-name</i>> <monitor <i>host</i>> <mpls (ldp <i>FEC address</i> rsvp <i>label-switched-path-name</i>)> <no-resolve> <propagate-ttl> <routing-instance <i>routing-instance-name</i>> <source <i>source-address</i>> <tos <i>value</i>> <ttl <i>value</i>> <wait <i>seconds</i>></pre>
Syntax (QFX Series and OCX Series)	<pre>traceroute <i>host</i> <as-number-lookup> <bypass-routing> <gateway <i>address</i>> <inet> <inet6> <interface <i>interface-name</i>> <monitor <i>host</i>> <no-resolve> <routing-instance <i>routing-instance-name</i>> <source <i>source-address</i>> <tos <i>value</i>> <ttl <i>value</i>> <wait <i>seconds</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. mpls option introduced in Junos OS Release 9.2. Command introduced in Junos OS Release 11.1 for the QFX Series. propagate-ttl option introduced in Junos OS Release 12.1. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	Display the route that packets take to a specified network host. Use traceroute as a debugging tool to locate points of failure in a network.
Options	host —IP address or name of remote host.

as-number-lookup—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

bypass-routing—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

clns—(Optional) Trace the route belonging to the Connectionless Network Service (CLNS).

gateway address—(Optional) Address of a router or switch through which the route transits.

inet | inet6—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

interface *interface-name*—(Optional) Name of the interface over which to send packets.

logical-system *logical-system-name*—(Optional) Perform this operation on all logical systems or on a particular logical system.

monitor *host*—(Optional) Display real-time monitoring information for the specified host.

mpls (*ldp FEC address* | *rsvp label-switched-path name*)—(Optional) See *traceroute mpls ldp* and *traceroute mpls rsvp*.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

propagate-ttl—(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-propagate-ttl** configuration statement.



NOTE: Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the traceroute attempt.

source *source-address*—(Optional) Source address of the outgoing traceroute packets.

tos *value*—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is 0 through 255.

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is 0 through 128.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

Required Privilege Level network

Related Documentation

- *traceroute monitor*

List of Sample Output

[traceroute on page 476](#)
[traceroute as-number-lookup host on page 476](#)
[traceroute no-resolve on page 477](#)
[traceroute propagate-ttl on page 477](#)
[traceroute \(Between CE Routers, Layer 3 VPN\) on page 477](#)
[traceroute \(Through an MPLS LSP\) on page 477](#)

Output Fields Table 120 on page 476 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 120: traceroute Output Fields

Field Name	Field Description
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.
<i>number-of-hops</i>	Number of hops from the source to the named router or switch.
<i>router-name</i>	Name of the router or switch for this hop.
<i>address</i>	Address of the router or switch for this hop.
Round trip time	Average round-trip time, in milliseconds (ms).

Sample Output

traceroute

```
user@host> traceroute santacruz
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1 blue23 (10.168.1.254)  2.370 ms  2.853 ms  0.367 ms
 2 red14 (10.168.255.250) 0.778 ms  2.937 ms  0.446 ms
 3 yellow (10.156.169.254) 7.737 ms  89.905 ms  0.834 ms
```

traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1 10.39.1.1 (10.39.1.1) 0.779 ms 0.728 ms 0.562 ms
 2 10.39.1.6 (10.39.1.6) [AS 32] 0.657 ms 0.611 ms 0.617 ms
 3 10.100.1.1 (10.100.1.1) [AS 10, 40, 50] 0.880 ms 0.808 ms 0.774 ms
```

traceroute no-resolve

```

user@host> traceroute santacruz no-resolve
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms

```

traceroute propagate-ttl

```

user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
    MPLS Label=299776 CoS=0 TTL=1 S=0
    MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
    MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms

```

traceroute (Between CE Routers, Layer 3 VPN)

```

user@host> traceroute vpn09
traceroute to host2.example.com (10.255.14.179), 30 hops max, 40
byte packets
 1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
 2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
    MPLS Label=100006 CoS=0 TTL=1 S=1
 3  host2.example.com (10.255.14.179)  0.783 ms  0.716 ms  0.686

```

traceroute (Through an MPLS LSP)

```

user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
    MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms

```


PART 6

Index

- [Index on page 481](#)

Index

Symbols

#, comments in configuration statements.....	xxii
(), in syntax descriptions.....	xxii
< >, in syntax descriptions.....	xxii
[], in configuration statements.....	xxii
{ }, in configuration statements.....	xxii
(pipe) command.....	3
(pipe), in syntax descriptions.....	xxii

A

accounting options	
configuration.....	11
overview.....	7
accounting profiles	
filter.....	20
interface.....	17
MIB.....	30
Routing Engine.....	32
Accounting-Options Configuration Statement	
Hierarchy.....	310
accounting-options statement.....	314
action-profile statement.....	315
adaptive services interfaces	
alarm conditions and configuration	
options.....	39
alarm class See alarm severity	
ALARM LED, color.....	37
alarm severity	
configuring for an interface.....	43
major (red)	38
See also major alarms	
minor (yellow).....	38
See also minor alarms	
alarms.....	450
active, displaying at login.....	43
conditions, on an interface.....	39
configurable.....	39
configuration requirements for interface	
alarms.....	43
licenses.....	42
major See major alarms	

minor See minor alarms	
overview.....	37
red See major alarms	
rescue configuration.....	42
severity See alarm severity	
types.....	37
verifying.....	45
yellow See minor alarms	
alias, CoS value.....	129
archive-sites statement	
accounting.....	316
usage guidelines.....	17
arithmetic operators, for multicast traffic.....	284
AS path, displaying.....	169

B

BGP (Border Gateway Protocol)	
monitoring.....	173
peers, probes to See BGP RPM probes	
RPM probes to BGP neighbors See BGP RPM	
probes	
statistics.....	174
BGP groups, displaying.....	174
BGP neighbors	
directing RPM probes to.....	65
displaying.....	174
monitoring with RPM probes.....	62
BGP peers See BGP neighbors	
BGP routing information.....	173
BGP RPM probes	
directing to select BGP neighbors	
(configuration editor).....	65
overview.....	53
setting up on local and remote device	
(configuration editor).....	62
BGP sessions, status.....	174
binary operators, for multicast traffic.....	284
braces, in configuration statements.....	xxii
brackets	
angle, in syntax descriptions.....	xxii
square, in configuration statements.....	xxii
built-in Ethernet ports See Ethernet ports;	
management interfaces	
bypass LSPs, testing.....	393
C	
capture-file statement.....	317
capturing packets See packet capture	

chassis	
monitoring.....	198
power management.....	198
chassis clusters	
redundancy group IP address monitoring	
configuration example.....	87
class-usage-profile statement.....	318
usage guidelines.....	28
classifiers, CoS.....	128, 135
clear chassis cluster ip-monitoring	
failure-count.....	371
clear chassis cluster ip-monitoring failure-count	
ip-address.....	372
CLI configuration editor	
interface alarms.....	43
RPM.....	55
cluster statement.....	319
code point aliases, CoS.....	129
comments, in configuration statements.....	xxii
configuration	
displaying	
current configuration.....	403
connections	
testing	
MPLS Layer 2 circuit connections.....	379
MPLS Layer 2 VPN connections.....	382
MPLS Layer 3 VPN connections.....	385
MPLS LDP connections.....	388
MPLS LSP-endpoint connections.....	391
MPLS RSVP connections.....	393
Content Filtering	
verifying.....	207
conventions	
text and syntax.....	xxi
CoS (class of service)	
classifiers.....	128, 135
CoS value aliases.....	129
forwarding classes.....	131
interfaces.....	127
loss priority.....	133
packet loss priority.....	133
RED drop profiles.....	130
rewrite rules.....	132
RPM probe classification.....	59
See also TCP RPM probes; UDP RPM probes	
scheduler maps.....	133
CoS components for link services	
applying on constituent links.....	293
counters statement.....	320
curly braces, in configuration statements.....	xxii
customer support.....	xxiii
contacting JTAC.....	xxiii
D	
datapath-debug	
security.....	321
datapath-debug statement.....	321
decryption-failures statement.....	322
destination-classes statement.....	323
usage guidelines.....	28
destination-interface statement	
RPM.....	324
destination-port statement	
RPM.....	325
device	
packet capture.....	266
diagnosis	
alarm configurations.....	45
CLI command summary.....	5
displaying firewall filter for.....	275
displaying packet capture configurations.....	270
interfaces.....	39, 137
J-Web tools overview.....	4
license infringement.....	42
load balancing on the link services	
interface.....	299
monitoring network performance.....	50
MPLS connections (J-Web).....	247
network traffic.....	281
packet capture.....	266
packet capture (J-Web).....	285
packet encapsulation on link services	
interfaces.....	298
ping command.....	250
ping host (J-Web).....	252
ping MPLS (J-Web).....	247
ports.....	39
preparation.....	249
system operation.....	239
traceroute (J-Web).....	241
traffic analysis with packet capture.....	266
verifying captured packets.....	270
verifying RPM probe servers.....	62
verifying RPM statistics.....	58
diagnostic commands.....	5
DiffServ code points, bits for RPM probes.....	68

- disabling
 - packet capture.....278
- DNS name resolution
 - troubleshooting.....292
- documentation
 - comments on.....xxiii
- drop probabilities, CoS.....130
- drop profiles, CoS.....130
- DS1 ports *See* T1 ports
- DS3 ports *See* E3 ports; T3 ports
- DSCPs (DiffServ code points), bits for RPM
 - probes.....68
- E**
 - E3 ports, alarm conditions and configuration
 - options.....39
 - egress *See* RPM probes, outbound times
 - encapsulation overhead, PPP and MLPPP.....298
 - encapsulation type
 - verifying for LFI and load balancing.....298
 - encapsulation, modifying on packet
 - capture-enabled interfaces.....280
 - Ethernet ports
 - alarm conditions and configuration
 - options.....39
 - configuring alarms on.....43
 - event viewer, J-Web
 - overview.....193
 - See also* system log messages
 - example
 - IP Monitoring.....79
- F**
 - FAQ (frequently asked questions)
 - Are LFI and load balancing working
 - correctly?.....295
 - What causes jitter and latency on multilink
 - bundles?.....295
 - Which CoS components apply on link services
 - interface?.....293
 - fields statement
 - for interface profiles.....326
 - usage guidelines.....18
 - for Routing Engine profiles.....327
 - usage guidelines.....32
 - file management
 - packet capture file creation.....267
 - file statement
 - accounting (associating with profile).....328
 - usage guidelines (filter profile).....21
 - usage guidelines (interface profile).....18
 - usage guidelines (MIB profile).....30
 - usage guidelines (Routing Engine
 - profile).....33
 - accounting (configuring log file).....329
 - usage guidelines.....14
 - files
 - status of, displaying.....373
 - files statement.....330
 - filter profile.....20
 - filter-profile statement.....331
 - usage guidelines.....20
 - filtering
 - command output.....3
 - firewall filters
 - for packet capture, configuring.....273
 - for packet capture, overview.....267
 - flow statement
 - (Security Flow).....332
 - font conventions.....xxi
 - forwarding classes, CoS.....131
 - fragmentation, verifying on the link services
 - interface.....297
 - frequency, test *See* RPM probes, test intervals
- G**
 - global-threshold statement.....334
 - global-weight statement.....335
 - groups
 - BGP, displaying.....174
- H**
 - hardware
 - major (red) alarm conditions on.....38
 - timestamp *See* RPM probe timestamps
 - hardware-timestamp statement.....335
 - heat status, checking.....198
 - host reachability
 - ping command.....250
 - ping host (J-Web).....252
 - hostname
 - monitoring traffic by matching.....283
 - pinging (CLI).....250
 - pinging (J-Web).....252
 - tracing a route to (CLI).....98, 245
 - tracing a route to (J-Web).....241

hosts, reachability	
MPLS Layer 2 circuits.....	379
MPLS Layer 2 VPN connections.....	382
MPLS Layer 3 VPN connections.....	385
MPLS LDP LSPs.....	388
MPLS LSP endpoints.....	391
MPLS RSVP LSPs.....	393
HTTP (Hypertext Transfer Protocol), RPM	
probes.....	50
Hypertext Transfer Protocol, RPM probes.....	50
I	
ICMP (Internet Control Message Protocol)	
RPM probes, description.....	50
RPM probes, inbound and outbound times.....	51
RPM probes, setting.....	55
icmp statement	
RPM.....	336
ICMPv6 (Internet Control Message Protocol)	
RPM probes, setting.....	66
idp potential violation statement.....	336
inbound time <i>See</i> RPM probes	
inet6-options statement	
RPM.....	337
ingress <i>See</i> RPM probes, inbound times	
Instance to which this connection belongs	
description.....	248
using.....	255
interface profile.....	17
interface-profile statement.....	338
usage guidelines.....	17
interfaces <i>See</i> management interfaces; network	
interfaces; ports	
interval statement	
accounting.....	339
usage guidelines (filter profile).....	21
usage guidelines (interface profile).....	19
usage guidelines (MIB profile).....	31
usage guidelines (Routing Engine	
profile).....	33
intervals, probe and test <i>See</i> RPM probes	
IP Monitoring.....	77, 79
route failover.....	77
supported threshold.....	78
test parameter.....	78
IP multicast	
tracing routes	
listen for responses.....	377
ip-monitoring statement.....	340
J	
J-Web configuration editor	
interface alarms.....	43
RPM.....	55
J-Web interface	
Diagnose options.....	5
event viewer.....	193
jitter	
description.....	51
<i>See also</i> RPM probes	
in RPM probes, improving with	
timestamps.....	50
monitoring.....	72
threshold, setting.....	68
jitter, removing on multilink bundles.....	295
Junos OS CLI	
diagnostic command summary.....	5
filtering command output.....	3
L	
label-switched paths <i>See</i> LSPs	
laptop <i>See</i> management device	
latency, in RPM probes, improving with	
timestamps.....	50
latency, reducing on multilink bundles.....	295
Layer 2 circuits	
reachability, testing.....	379
Layer 2 circuits, monitoring.....	247
Layer 2 VPNs	
reachability, testing.....	382
Layer 2 VPNs, monitoring.....	247
Layer 3 VPNs	
reachability, testing.....	385
Layer 3 VPNs, monitoring.....	247
LDP LSPs	
ping interval.....	388
libpcap format, for packet capture files.....	270
licenses, alarm conditions and remedies.....	42
limitations	
ALARM LED lights yellow whether alarm is	
minor or major.....	37
MPLS, no LSP statistics on outbound	
device.....	145
mtrace from-source packet statistics always	
0.....	95
performance degradation with monitor traffic	
command.....	281
PPP, no J-Web monitoring information	
available.....	149

- link services interface
 - applying CoS components on constituent links.....293
 - fragmentation, troubleshooting.....297
 - load balancing, troubleshooting.....299
 - MLPPP header overhead.....298
 - packet encapsulation, troubleshooting.....298
 - PPP header overhead.....298
 - preventing dropped packets on PVCs.....302
 - reducing jitter and latency on multilink bundles.....295
 - troubleshooting LFI and load balancing.....295
- load balancing on link services interfaces
 - FAQ.....295
 - troubleshooting.....295
 - verifying.....299
- Locate LSP from interface name
 - description.....248
 - using.....255
- Locate LSP from virtual circuit information
 - description.....248
 - using.....255
- Locate LSP using interface name
 - description.....248
 - using.....255
- log files
 - display of
 - starting.....374
 - stopping.....376
 - status, displaying.....373
- logical interfaces, CoS.....127
- logical operators, for multicast traffic.....284
- loss priority, CoS.....133
- LSPs
 - LDP, ping interval.....388
 - RSVP, ping interval.....393
- LSPs (label-switched paths)
 - information about.....144
 - monitoring, with ping MPLS.....247
 - statistics.....146
- M**
- major (red) alarms
 - description.....38
- management device
 - diagnosing problems from.....4
 - monitoring from.....3
 - recovering root password from.....291
- management interfaces
 - alarm conditions and configuration
 - options.....39
 - configuring alarms on.....43
 - monitoring.....137, 142
 - statistics.....137
- manuals
 - comments on.....xxiii
- match conditions, for multicast traffic
 -283
- maximum-capture-size
 - security log.....342
- MIB profile.....30
- mib-profile statement.....343
 - usage guidelines.....30
- minimum accounting options configuration.....13
- minor (yellow) alarms
 - description.....38
- MLPPP encapsulation, on the link services interface.....298
- monitor interface command.....137
 - controlling output.....138
- monitor interface traffic command.....137
 - controlling output.....138
- monitor list command.....239, 373
- monitor start command.....239, 374
- monitor stop command.....239, 376
- monitor traffic command.....281
 - options.....281
 - performance impact.....281
- monitor traffic matching command.....281
 - arithmetic, binary, and relational operators.....284
 - logical operators.....284
 - match conditions.....283
- monitoring
 - BGP.....174
 - BGP neighbors, with RPM probes.....62
 - chassis.....198
 - interfaces.....137, 142
 - Layer 2 circuits.....247
 - Layer 2 VPNs.....247
 - Layer 3 VPNs.....247
 - MPLS traffic
 - engineering.....143, 144, 145, 146, 148
 - network interface traffic.....281
 - network traffic with packet capture.....266
 - OSPF.....172
 - ports.....142

PPP (CLI).....	149	network interfaces	
PPPoE.....	149	alarm conditions and configuration	
preparation.....	249	options.....	39
RIP.....	170	configuring alarms on.....	43
routing information.....	168	integrated services, alarm conditions and	
routing tables.....	168	configuration options.....	39
RPM probes.....	72	monitoring.....	137, 142
system logs.....	239	monitoring MPLS traffic engineering.....	144
trace files.....	239	monitoring traffic.....	281
monitoring the wx interface.....	154	monitoring, CoS.....	127
MPLS		monitoring, PPPoE.....	150
Layer 2 circuit connections		monitoring, RSVP.....	148
operability, checking.....	379	packet capture, configuring on.....	271
Layer 2 VPN connections		packet capture, disabling before changing	
operability, checking.....	382	encapsulation.....	280
Layer 3 VPN connections		packet capture, supported on.....	266
operability, checking.....	385	services, alarm conditions and configuration	
LDP-signaled LSP connections		options.....	39
operability, checking.....	388	statistics.....	137
LSP endpoint connections		network performance See RPM	
operability, checking.....	391	next hop, displaying.....	169
MPLS (Multiprotocol Label Switching)		nonpersistent statement.....	345
connections, checking.....	247	accounting	
LSPs.....	144	usage guidelines.....	15
monitoring interfaces.....	144	O	
monitoring LSP information.....	144	object-names statement.....	345
monitoring LSP statistics.....	145, 146	objects-names statement	
monitoring MPLS interfaces.....	144	for Routing Engine profiles	
monitoring RSVP interfaces.....	148	usage guidelines.....	31
monitoring RSVP sessions.....	146	Open Shortest Path First See OSPF	
monitoring traffic engineering.....	143	operation statement.....	346
mpls statement.....	344	for MIB profiles	
mtrace monitor command.....	240, 377	usage guidelines.....	31
results.....	240	operational mode, filtering command output.....	3
mtrace-from-source command.....	95	operators	
options.....	95	arithmetic, binary, and relational	
results.....	97	operators.....	284
multicast		logical.....	284
trace operations, displaying.....	240	OSPF (Open Shortest Path First)	
tracing paths.....	95	monitoring.....	171
multilink bundles		statistics.....	172
preventing dropped packets.....	302	OSPF interfaces	
reducing latency.....	295	displaying.....	172
removing jitter.....	295	status.....	172
Multiprotocol Label Switching See MPLS		OSPF neighbors	
N		displaying.....	172
neighbors, BGP See BGP neighbors; BGP RPM		status.....	172
probes		OSPF routing information.....	171

outbound time See RPM probes

P

- P2MP LSPs, testing.....393
- packet capture
 - configuring.....271
 - configuring (J-Web).....285
 - configuring on an interface.....271
 - device interfaces supported.....266
 - disabling.....278
 - disabling before changing encapsulation on
 - interfaces.....280
 - displaying configurations.....270
 - displaying firewall filter for.....275
 - enabling.....268
 - encapsulation on interfaces, disabling before
 - modifying.....280
 - files See packet capture files
 - firewall filters, configuring.....273
 - firewall filters, overview.....267
 - J-Web tool.....285
 - overview.....266
 - overview (J-Web).....285
 - preparation.....268
 - verifying captured packets.....270
 - verifying configuration.....270
 - verifying firewall filter for.....275
- packet capture configuration
 - datapath debugging.....275
- packet capture files
 - analyzing.....267
 - libpcap format.....270
 - overview.....267
 - renaming before modifying encapsulation on
 - interfaces.....280
- Packet Capture page
 - field summary.....286
 - results.....288
- packet encapsulation
 - troubleshooting on the link services
 - interface.....295
 - verifying on the link services interface.....298
- packet fragmentation
 - troubleshooting on the link services
 - interface.....295
 - verifying on the link services interface.....297
- packet loss priority, CoS.....133
- packet-capture statement.....347
- packet-filter statement
 - security.....348
- packets
 - capturing.....266
 - capturing with J-Web packet capture.....285
 - monitoring jitter.....72
 - monitoring packet loss.....72
 - monitoring round-trip times.....72
 - multicast, tracking95
 - packet capture.....266
 - packet capture (J-Web).....285
 - tracking MPLS.....258
 - tracking with J-Web traceroute.....241
- parentheses, in syntax descriptions.....xxii
- passwords
 - srx root password, recovering.....291
- PC See management device
- PCAP See packet capture
- peers, BGP See BGP neighbors; BGP RPM probes
- performance, monitoring See RPM
- physical interfaces, CoS.....127
- PIMs (Physical Interface Modules)
 - checking power and heat status.....198
- ping
 - host reachability (CLI).....250
 - host reachability (J-Web).....252
 - ICMP probes.....55
 - RPM probes See RPM probes
 - TCP and UDP probes.....59
- ping command.....250
 - options.....250
- Ping end point of LSP
 - description.....248
 - using.....255
- Ping Host page
 - field summary.....252
- Ping LDP-signaled LSP
 - description.....248
 - using.....255
- Ping LSP to Layer 3 VPN prefix
 - description.....248
 - using.....255
- ping MPLS (J-Web)
 - indications.....258
 - Layer 2 circuits.....247
 - Layer 2 VPNs.....247
 - Layer 3 VPNs.....247
 - LSP state.....247
 - options.....247, 248

requirements.....	249	probe-limit statement.....	351
results.....	258	probe-server statement.....	352
ping mpls l2circuit command.....	259, 379	probe-type statement.....	353
results.....	258	probes, monitoring.....	72, 149
ping mpls l2vpn command.....	260, 382	<i>See also</i> RPM probes	
results.....	258	profiles, accounting	
ping mpls l3vpn command.....	261, 385	filter.....	20
results.....	258	interface.....	17
ping mpls ldp command.....	262, 388	MIB.....	30
results.....	258	Routing Engine.....	32
ping mpls lsp-end-point command.....	262, 391	protocols	
results.....	258	originating, displaying.....	169
Ping MPLS page		OSPF, monitoring.....	171
field summary.....	255	PPP, monitoring.....	149
results.....	258	RIP, monitoring.....	170
ping mpls rsvp command.....	262, 393	routing protocols, monitoring.....	168, 173
results.....	258	PVCs (permanent virtual circuits)	
Ping RSVP-signaled LSP		preventing dropped packets on.....	302
description.....	248		
using.....	255	Q	
pipe () command, to filter output.....	3	Quick Configuration	
Point-to-Point Protocol <i>See</i> PPP		RPM pages.....	55
Point-to-Point Protocol over Ethernet <i>See</i> PPPoE			
ports		R	
alarm conditions and configuration		random early detection (RED) drop profiles,	
options.....	39	CoS.....	130
configuring alarms on.....	43	real-time monitoring	
individual port types.....	39	files.....	373
monitoring.....	142	real-time performance monitoring <i>See</i> RPM	
power management, chassis.....	198	RED drop profiles, CoS.....	130
PPP (Point-to-Point Protocol)		redundancy-group statement.....	354
monitoring (CLI).....	149	redundant Ethernet interface LAG.....	81
PPP encapsulation		relational operators, for multicast traffic.....	284
on the link services interface.....	298	request ppoe connect command.....	398
PPPoE (Point-to-Point Protocol over Ethernet)		request ppoe disconnect command.....	399
interfaces.....	150	rescue configuration, alarm about.....	42
monitoring.....	149	Resource Reservation Protocol <i>See</i> RSVP	
session status.....	150	reth	
statistics.....	150	link aggregation group.....	81
version information.....	150	retry-interval statement.....	355
Primary-level entry		rewrite rules, CoS.....	132
secondary-level entry.....	233	RIP (Routing Information Protocol)	
Primary-level entry only.....	233	monitoring.....	170
probe loss		statistics.....	170
monitoring.....	72	RIP neighbors	
threshold, setting.....	68	displaying.....	170
probe statement		status.....	170
RPM.....	349	RIP routing information.....	170
probe-interval statement.....	350	root password recovery.....	291

round-trip time		RPM pages.....	55
description.....	51	field summary.....	68
<i>See also</i> RPM probes		RPM probe timestamps	
threshold, setting.....	68	overview.....	50
routes, displaying		setting (configuration editor).....	55
to specified network host.....	474	RPM probes	
routing		basic (configuration editor).....	55
monitoring.....	168	BGP neighbors <i>See</i> BGP RPM probes	
traceroute (J-Web).....	241	cumulative jitter.....	72
Routing Engine profile.....	32	current tests.....	72
routing solutions		DSCP bits (Quick Configuration).....	68
applying CoS components on link services		graph results.....	72
interface.....	293	ICMP (configuration editor).....	55
load balancing on link services interfaces.....	295	ICMPv6 (configuration editor).....	66
preventing dropped packets on PVCs.....	302	inbound times.....	51
reducing jitter and latency on multilink		IPv6 (configuration editor).....	66
bundles.....	295	jitter threshold.....	68
routing table		monitoring.....	72
monitoring.....	168	outbound times.....	51
routing-engine-profile statement.....	356	probe count, setting (Quick	
usage guidelines.....	32	Configuration).....	68
RPM (real-time performance monitoring)		probe count, tuning.....	67
basic probes (configuration editor).....	55	probe counts.....	51
BGP monitoring <i>See</i> BGP RPM probes		probe intervals.....	50
inbound and outbound times.....	51	probe intervals, setting (Quick	
IPv6 probes (configuration editor).....	66	Configuration).....	68
jitter, viewing.....	72	probe intervals, tuning.....	67
monitoring probes.....	72	probe loss count.....	68
overview.....	50	probe owner.....	68
<i>See also</i> RPM probes		probe type, setting (Quick Configuration).....	68
preparation.....	55	probe types.....	50
probe and test intervals.....	50	round-trip time threshold.....	68
probe counts.....	51	round-trip times, description.....	51
Quick Configuration.....	55	round-trip times, viewing.....	72
round-trip times, description.....	51	SNMP traps (Quick Configuration).....	68
round-trip times, viewing.....	72	source address, setting.....	67
sample configuration.....	58	TCP (configuration editor).....	59
sample graphs.....	72	<i>See also</i> TCP RPM probes	
statistics.....	51	TCP server port.....	68
statistics, verifying.....	58	test intervals.....	50
TCP probes (configuration editor).....	59	test intervals, setting (Quick	
<i>See also</i> TCP RPM probes		Configuration).....	68
tests.....	50	test target.....	68
tests, viewing.....	72	threshold values, description.....	53
threshold values.....	53	threshold values, setting (Quick	
tuning probes.....	67	Configuration).....	68
UDP probes (configuration editor).....	59	timestamps <i>See</i> RPM probe timestamps	
<i>See also</i> UDP RPM probes		tuning.....	67
verifying probe servers.....	62		

UDP (configuration editor).....	59
See also UDP RPM probes	
UDP server port.....	68
verifying TCP and UDP probe servers.....	62
RSVP	
LSP connections	
operability, checking.....	393
RSVP (Resource Reservation Protocol)	
interfaces, monitoring.....	148
sessions, monitoring.....	146
RSVP LSPs	
ping interval.....	393
RTT See RPM probes, round-trip times	

S

samples

alarm configuration.....	45
basic RPM probes.....	55
RPM probes.....	58
RPM test graphs.....	72
TCP and UDP probes.....	59

scheduler maps, CoS.....	133
--------------------------	-----

security

alarms.....	450
packet capture for intrusion detection.....	266

Security Configuration Statement Hierarchy.....	358
---	-----

security policy

DNS name resolution.....	292
--------------------------	-----

serial ports

alarm conditions and configuration	
options.....	39
configuring alarms on.....	43

services module

alarm conditions and configuration	
options.....	39

Services Router

monitoring	3
performance monitoring.....	50

sessions

BGP peer, status details.....	174
RSVP, monitoring.....	146

severity levels

for alarms See alarm severity

show bgp neighbor command.....	173
show bgp summary command.....	173
show chassis alarms command.....	45, 46, 401
show chassis cluster ip-monitoring status	
redundancy-group command.....	406
show chassis environment command.....	198

show chassis hardware command.....	196, 198
show chassis power-ratings command.....	198
show chassis redundant-power-supply	
command.....	198
show chassis routing-engine command.....	198
show class-of-service classifier	
command.....	128, 135
show class-of-service code-point-aliases	
command.....	129
show class-of-service drop-profile command.....	130
show class-of-service forwarding-class	
command.....	131
show class-of-service rewrite-rules command.....	132
show class-of-service scheduler-map	
command.....	133
show configuration command.....	403
show firewall filter dest-all command.....	275
show interfaces command.....	409
show interfaces detail command.....	142
show interfaces interface-name command.....	142
show interfaces pp0 command.....	149
show interfaces terse command.....	142
show mpls interface command.....	144
show mpls lsp command.....	144
show mpls statistics command.....	145
show ospf interfaces command.....	171
show ospf neighbors command.....	171
show ospf statistics command.....	171
show poe interface command.....	440
show poe telemetries.....	442
show ppp address-pool command.....	149
show ppp interface command.....	149
show ppp statistics command.....	149
show ppp summary command.....	149
show pppoe interfaces command.....	149, 444
show pppoe statistics command.....	149, 448
show pppoe version command.....	149
show redundant-power-supply command.....	198
show rip neighbors command.....	170
show rip statistics command.....	170
show route detail command.....	168
show route terse command.....	168
show security alarms command.....	450
show security datapath-debug capture.....	454
show security datapath-debug counter.....	455
show security monitoring fpc fpc-number	
command.....	458
show services ip-monitoring status	
command.....	464

- show services rpm active-servers
 - command.....62, 64
 - explanation.....62, 64
- show services rpm probe-results
 - command.....58, 72, 468
 - explanation.....58
- show system alarms command.....473
- show system storage command.....196
- show system uptime command.....196
- show system users command.....196
- show version command.....196
- show forwarding-options command.....270
- size statement
 - accounting.....360
 - usage guidelines.....16
- SNMP traps
 - performance monitoring See RPM probes
- source-classes statement.....360
 - usage guidelines.....28
- SRX Series
 - alarms.....37
 - monitoring3
 - packet capture.....266
 - performance monitoring.....50
- start-time statement
 - accounting.....361
 - usage guidelines.....16
- statistics
 - BGP174
 - interfaces.....137
 - LSP.....146
 - OSPF.....172
 - performance monitoring.....51
 - PPPoE.....150
 - RIP.....170
 - RPM, description.....51
 - RPM, monitoring.....72
 - RPM, verifying.....58
- status
 - BGP174
 - OSPF interfaces.....172
 - OSPF neighbors.....172
 - RIP neighbors.....170
- support, technical See technical support
- syntax conventions.....xxi
- system log messages
 - event viewer.....193
- system logs
 - monitoring.....239
- system management
 - displaying log and trace file contents.....239
- T**
 - T1 ports
 - alarm conditions and configuration
 - options.....39
 - configuring alarms on.....43
 - T3 ports
 - alarm conditions and configuration
 - options.....39
 - configuring alarms on.....43
 - target statement.....362
 - RPM.....362
 - TCP RPM probes
 - CoS classification, destination interface
 - requirement.....59
 - CoS classification, use with caution.....59
 - description.....50
 - server port.....68
 - setting.....59
 - verifying servers.....62
 - technical support
 - contacting JTAC.....xxiii
 - temporary files
 - for packet capture.....267
 - tests See RPM
 - threshold values, for RPM probes See RPM probes
 - thresholds statement
 - RPM.....363
 - timestamps
 - for RPM probes See RPM probe timestamps
 - suppressing in packet headers, in captured
 - packets.....286
 - suppressing in packet headers, in traffic
 - monitoring.....282
 - trace files
 - display of
 - starting.....374
 - stopping.....376
 - monitoring.....239
 - multicast, monitoring.....240
 - status, displaying.....373
 - traceoptions statement
 - datapath-debug.....365
 - traceroute
 - CLI command.....245
 - indications.....243
 - J-Web tool.....241

results.....	243
TTL increments.....	241
tracert command.....	245, 474
options.....	245
tracert monitor	
CLI command.....	97
tracert monitor command.....	97
options.....	98
results.....	99
Tracert page	
field summary.....	241
tracing routes	
monitoring.....	377
traffic	
analyzing with packet capture.....	266
multicast, tracking.....	95
tracking with J-Web tracert.....	241
transfer-interval statement	
accounting.....	366
usage guidelines.....	16
traps statement.....	367
troubleshooting	
applying CoS components on link services	
interface.....	293
DNS name resolution in security policy.....	292
dropped packets on PVCs.....	302
jitter and latency on multilink bundles.....	295
LFI and load balancing on multilink	
bundles.....	295
packet capture for analysis.....	266
<i>See also</i> diagnosis; packet capture	
root password recovery.....	291
TTL (time to live)	
default, in multicast path-tracking queries.....	95
increments, in tracert packets.....	241
threshold, in multicast trace results.....	97
total, in multicast trace results.....	97

U

UDP RPM probes	
CoS classification, destination interface	
requirement.....	59
CoS classification, use with caution.....	59
description.....	50
server port.....	68
setting.....	59
verifying servers.....	62

V

verification	
alarm configurations.....	45
captured packets.....	270
destination path (J-Web).....	241
firewall filter for packet capture.....	275
host reachability (CLI).....	250
host reachability (J-Web).....	252
load balancing on the link services	
interface.....	299
LSPs (J-Web).....	247
packet capture.....	270
packet encapsulation on link services	
interface.....	298
RPM configuration.....	58
RPM probe servers.....	62, 64
RPM statistics.....	58
tracing multicast paths.....	95
version	
PPPoE, information about.....	150
View Events page	
field summary (filtering log messages).....	169

W

Web Filtering	
verifying.....	214

Y

yellow alarms *See* minor alarms