



Junos[®] OS for EX Series Ethernet Switches

Network Management and Monitoring on EX Series Switches

Release

14.1X53



Published: 2014-12-18

Juniper Networks, Inc.
1194 North Mathilda Avenue
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 for EX Series Ethernet Switches Network Management and Monitoring on EX Series Switches
Release 14.1X53
Copyright © 2014, 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	xvii
	Documentation and Release Notes	xvii
	Supported Platforms	xvii
	Using the Examples in This Manual	xvii
	Merging a Full Example	xviii
	Merging a Snippet	xviii
	Documentation Conventions	xix
	Documentation Feedback	xxi
	Requesting Technical Support	xxi
	Self-Help Online Tools and Resources	xxi
	Opening a Case with JTAC	xxii
Part 1	Overview	
Chapter 1	Real-Time Performance Monitoring	3
	Understanding Real-Time Performance Monitoring on EX Series Switches	3
	RPM Packet Collection	4
	Tests and Probe Types	4
	Hardware Timestamps	4
	Limitations of RPM on EX Series Switches	6
Chapter 2	Port Mirroring	9
	Understanding Port Mirroring on EX Series Switches	9
	Port Mirroring Overview	9
	Port Mirroring Terminology	10
	Configuration Guidelines for Port Mirroring on the Switches	12
Chapter 3	sFlow Monitoring Technology	17
	Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch	17
	Sampling Mechanism and Architecture of sFlow Technology on EX Series Switches	17
	Adaptive Sampling	18
	sFlow Agent Address Assignment	19
Chapter 4	Ethernet OAM Link Fault Management and Connectivity Fault Management	21
	Understanding Ethernet OAM Link Fault Management for an EX Series Switch	21
	Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch	22

	Understanding Ethernet Frame Delay Measurements on Switches	24
	Ethernet Frame Delay Measurements	24
	Types of Ethernet Frame Delay Measurements	25
	Limitations	26
Chapter 5	Uplink Failure Detection	27
	Understanding Uplink Failure Detection	27
	Uplink Failure Detection Overview	27
	Failure Detection Pair	28
Part 2	Configuration	
Chapter 6	Configuration Examples	33
	Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches	33
	Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches	37
	Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches	42
	Example: Configuring Port Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX Series Switches	51
	Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches	61
	Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches	68
	Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches	72
Chapter 7	Configuration Tasks	75
	Configuring SNMP (J-Web Procedure)	75
	Configuring Port Mirroring to Analyze Traffic (CLI Procedure)	78
	Configuring Port Mirroring for Local Traffic Analysis	79
	Configuring Port Mirroring for Remote Traffic Analysis	79
	Filtering the Traffic Entering an Analyzer	80
	Configuring Port Mirroring to Analyze Traffic (J-Web Procedure)	81
	Configuring sFlow Technology for Network Monitoring (CLI Procedure)	84
	Configuring Real-Time Performance Monitoring (J-Web Procedure)	85
	Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch (CLI Procedure)	92
	Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)	93
	Creating the Maintenance Domain	93
	Configuring the Maintenance Domain MIP Half Function	94
	Creating a Maintenance Association	94
	Configuring the Continuity Check Protocol	95
	Configuring a Maintenance Association End Point	95
	Configuring a Connectivity Fault Management Action Profile	96
	Configuring the Linktrace Protocol	96
	Configuring Ethernet OAM Link Fault Management (CLI Procedure)	97
	Configuring Interfaces for Uplink Failure Detection (CLI Procedure)	99

	Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure)	100
	Configuring One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)	101
	Configuring an Iterator Profile on a Switch (CLI Procedure)	102
	Triggering an Ethernet Frame Delay Measurement Session on a Switch	103
	Configuring Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)	104
Chapter 8	Configuration Statements: SNMP	105
	[edit snmp] Configuration Statement Hierarchy on EX Series Switches	107
	Supported Statements in the [edit snmp] Hierarchy Level	108
	Unsupported Statements in the [edit snmp] Hierarchy Level	112
	address (SNMP)	114
	address-mask	115
	agent-address	115
	alarm (SNMP RMON)	116
	authorization	117
	bucket-size	117
	categories	118
	client-list	118
	client-list-name	119
	clients	119
	commit-delay	120
	community (SNMP)	121
	community (SNMP RMON)	122
	community-name (SNMP)	123
	contact (SNMP)	124
	description (SNMP)	125
	description (SNMP RMON)	126
	destination-port	126
	engine-id (SNMP)	127
	event (SNMP)	128
	falling-event-index	128
	falling-threshold	129
	falling-threshold	130
	falling-threshold-interval	131
	filter-duplicates	131
	filter-interfaces	132
	group (Configuring Access Privileges)	133
	group (Associating a Security Name)	134
	health-monitor	134
	history	135
	interface (SNMP RMON History)	136
	interface (SNMP)	136
	interval	137
	interval (SNMP Health Monitor)	137
	interval (SNMP RMON)	138
	location (SNMP)	138

logical-system (SNMP)	139
message-processing-model	140
name	140
nonvolatile	141
notify	142
notify-filter (Applying to the Management Target)	143
notify-filter (Configuring the Profile Name)	143
notify-view	144
oid (SNMP)	144
oid	145
owner	145
parameters	146
port (SNMP)	146
read-view	147
request-type	148
rising-event-index	149
rising-threshold (SNMP Health Monitor)	150
rising-threshold (SNMP RMON)	151
rmon	151
rmon	152
routing-instance (SNMP)	153
routing-instance (SNMPv3)	154
sample-type	154
security-level (Defining Access Privileges)	155
security-level (Generating SNMP Notifications)	156
security-model (Access Privileges)	157
security-model (Group)	158
security-model (SNMP Notifications)	159
security-name (Security Group)	160
security-name (Community String)	161
security-name (SNMP Notifications)	162
security-to-group	163
snmp	163
snmp	164
snmp-community	165
source-address (SNMP)	166
startup-alarm	167
syslog-subtag	167
tag (SNMPv3)	168
tag-list	168
target-address	169
target-parameters	170
targets	171
traceoptions (SNMP)	172
trap-group	174
trap-options	175
type (SNMP RMON)	176
type (SNMPv3)	177
v3	178

	vacm	180
	variable	181
	version (SNMP)	181
	view (Associating a MIB View with a Community)	182
	view (Configuring a MIB View)	183
	write-view	184
Chapter 9	Configuration Statements: Port Mirroring	185
	[edit ethernet-switching-options] Configuration Statement Hierarchy on EX Series Switches	185
	Supported Statements in the [edit ethernet-switching-options] Hierarchy Level	186
	Unsupported Statements in the [edit ethernet-switching-options] Hierarchy Level	188
	analyzer	189
	egress	190
	egress (Interface or VLAN)	190
	ethernet-switching-options	191
	ingress (vlans)	194
	ingress (Interface or VLAN)	195
	input (Port Mirroring)	196
	interface	197
	loss-priority (Port Mirroring)	198
	no-tag	199
	output (Port Mirroring)	200
	ratio	201
	vlan (Port Mirroring)	202
Chapter 10	Configuration Statements: sFlow Technology	203
	[edit protocols] Configuration Statement Hierarchy on EX Series Switches	203
	agent-id	205
	collector	205
	disable (sFlow Monitoring Technology)	206
	interfaces (sFlow Monitoring Technology)	207
	polling-interval	208
	sample-rate	209
	sflow	210
	source-ip	211
	udp-port	211
Chapter 11	Configuration Statements: Ethernet OAM Connectivity Fault Management	213
	action-profile (Applying to OAM CFM, for EX Series Switch Only)	215
	age (EX Series Switch Only)	216
	auto-discovery (EX Series Switch Only)	216
	calculation-weight	217
	connectivity-fault-management (EX Series Switch Only)	218
	continuity-check (EX Series Switch Only)	219
	cycle-time	220
	data-tlv-size	221

delay	222
delay-variation	223
direction (EX Series Switch Only)	224
hold-interval (OAM CFM, for EX Series Switch Only)	224
interface (OAM CFM, for EX Series Switch Only)	225
interval (EX Series Switch Only)	226
iteration-count	227
iteration-period	227
level (EX Series Switch Only)	228
linktrace (EX Series Switch Only)	228
loss-threshold (EX Series Switch Only)	229
maintenance-association (EX Series Switch Only)	230
maintenance-domain (EX Series Switch Only)	231
measurement-type (OAM LFM)	232
mep (EX Series Switch Only)	233
mip-half-function (EX Series Switch Only)	234
name-format (EX Series Switch Only)	235
path-database-size (EX Series Switch Only)	235
performance-monitoring (OAM LFM)	236
priority (Protocols OAM)	237
priority (OAM Connectivity-Fault Management)	237
remote-mep (EX Series Switch Only)	238
short-name-format	239
sla-iterator-profile	240
sla-iterator-profiles (OAM LFM)	241
traceoptions (OAM CFM, for EX Series Switch Only)	242
Chapter 12 Configuration Statements: Ethernet OAM Link Fault Management	245
action (OAM LFM)	246
action-profile	247
allow-remote-loopback	248
ethernet (OAM LFM)	249
event (OAM LFM)	252
event-thresholds	252
frame-error	253
frame-period	253
frame-period-summary	254
interface (OAM LFM)	255
link-adjacency-loss	256
link-discovery	256
link-down	257
link-event-rate	257
link-fault-management	258
negotiation-options	259
no-allow-link-events	259
oam	260
pdu-interval	262
pdu-threshold	263
remote-loopback	263

	symbol-period	264
	syslog (OAM LFM)	264
	traceoptions (OAM LFM)	265
Chapter 13	Configuration Statements: RPM	267
	data-fill	268
	data-size	269
	destination-port	270
	dscp-code-point	271
	hardware-timestamp	272
	history-size	273
	moving-average-size	273
	one-way-hardware-timestamp	274
	port (RPM)	274
	probe	275
	probe-count	276
	probe-interval	276
	probe-limit	277
	probe-server	278
	probe-type	279
	routing-instance	280
	routing-instances	280
	rpm (Interfaces)	281
	source-address (Services)	281
	tcp	282
	test	283
	test-interval	284
	thresholds	285
	traps	286
	udp	287
Chapter 14	Configuration Statements: Uplink Failure Detection	289
	action (Uplink Failure Detection)	289
	group (Uplink Failure Detection)	290
	link-to-disable	290
	link-to-monitor	291
	traceoptions (Uplink Failure Detection)	292
	uplink-failure-detection	293
Part 3	Administration	
Chapter 15	Routine Monitoring	297
	Monitoring Hosts Using the J-Web Ping Host Tool	297
	Monitoring Network Traffic Using Traceroute	299
	Verifying Input and Output for Port Mirroring Analyzers on EX Series Switches	300
	Viewing Real-Time Performance Monitoring Information	302
	Verifying That Uplink Failure Detection Is Working Correctly	302

Chapter 16	Operational Commands: General	305
	monitor traffic	306
	ping	316
	show pfe statistics bridge	320
	traceroute	325
Chapter 17	Operational Commands: RPM	329
	show services rpm active-servers	330
	show services rpm history-results	331
	show services rpm probe-results	334
Chapter 18	Operational Commands: SNMP	341
	clear snmp rmon history	342
	clear snmp statistics	343
	request snmp spoof-trap	345
	show snmp health-monitor	351
	show snmp inform-statistics	358
	show snmp mib	360
	show snmp rmon	363
	show snmp rmon history	367
	show snmp statistics	371
	show snmp v3	375
Chapter 19	Operational Commands: Port Mirroring	379
	show analyzer	380
Chapter 20	Operational Commands: sFlow	383
	show sflow	384
	show sflow interface	386
	show sflow collector	388
	clear sflow collectors statistics	389
Chapter 21	Operational Commands: Ethernet OAM Connectivity Fault Management	391
	clear oam ethernet connectivity-fault-management delay-statistics	392
	clear oam ethernet connectivity-fault-management sla-iterator-statistics	394
	clear oam ethernet connectivity-fault-management statistics	395
	monitor ethernet delay-measurement	396
	show oam ethernet connectivity-fault-management delay-statistics	401
	show oam ethernet connectivity-fault-management forwarding-state	405
	show oam ethernet connectivity-fault-management interfaces	409
	show oam ethernet connectivity-fault-management path-database	415
	show oam ethernet connectivity-fault-management mep-database	418
	show oam ethernet connectivity-fault-management mip	424
	show oam ethernet connectivity-fault-management sla-iterator-statistics	425
Chapter 22	Operational Commands: Ethernet OAM Link Fault Management	431
	show oam ethernet link-fault-management	432
Chapter 23	Operational Commands: Uplink Failure Detection	437
	show uplink-failure-detection	438

Part 4	Troubleshooting	
Chapter 24	Troubleshooting Procedures and Commands	443
	Troubleshooting Port Mirroring Configuration Error Messages	443
	An Analyzer Configuration Returns a "Multiple interfaces cannot be configured as a member of Analyzer output VLAN" Error Message . . .	443
	Troubleshooting sFlow on Management Interface	444
	show pfe statistics bridge	446

List of Figures

Part 1	Overview	
Chapter 1	Real-Time Performance Monitoring	3
	Figure 1: RPM Timestamps	5
Chapter 4	Ethernet OAM Link Fault Management and Connectivity Fault Management	21
	Figure 2: Relationship Among MEPs, MIPs, and Maintenance Domain Levels	24
Chapter 5	Uplink Failure Detection	27
	Figure 3: Uplink Failure Detection Configuration on Switches	28
Part 2	Configuration	
Chapter 6	Configuration Examples	33
	Figure 4: sFlow Technology Monitoring System	34
	Figure 5: Network Topology for Local Port Mirroring Example	39
	Figure 6: Remote Port Mirroring Example Network Topology	44
	Figure 7: Remote Port Mirroring Example Network Topology Using Multiple VLAN Member Interfaces	54
	Figure 8: Remote Port Mirroring Example Through a Transit Switch Network Topology	62

List of Tables

	About the Documentation	xvii
	Table 1: Notice Icons	xix
	Table 2: Text and Syntax Conventions	xix
Part 1	Overview	
Chapter 2	Port Mirroring	9
	Table 3: Port Mirroring Terminology	11
	Table 4: Configuration Guidelines for Port Mirroring	13
Part 2	Configuration	
Chapter 7	Configuration Tasks	75
	Table 5: SNMP Configuration Page	76
	Table 6: Port Mirroring Configuration Settings	82
	Table 7: RPM Probe Owner, Concurrent Probes, and Probe Servers Configuration Fields	87
	Table 8: Performance Probe Tests Configuration Fields	88
Chapter 8	Configuration Statements: SNMP	105
	Table 9: Unsupported [edit snmp] Configuration Statements on EX Series Switches	112
Part 3	Administration	
Chapter 15	Routine Monitoring	297
	Table 10: J-Web Ping Host Field Summary	298
	Table 11: Traceroute field summary	299
Chapter 16	Operational Commands: General	305
	Table 12: Match Conditions for the monitor traffic Command	308
	Table 13: Logical Operators for the monitor traffic Command	309
	Table 14: Arithmetic and Relational Operators for the monitor traffic Command	311
	Table 15: show pfe statistics bridge Output Fields	320
	Table 16: traceroute Output Fields	327
Chapter 17	Operational Commands: RPM	329
	Table 17: show services rpm active-servers Output Fields	330
	Table 18: show services rpm history-results Output Fields	331
	Table 19: show services rpm probe-results Output Fields	334
Chapter 18	Operational Commands: SNMP	341

	Table 20: show snmp health-monitor Output Fields	351
	Table 21: show snmp inform-statistics Output Fields	358
	Table 22: show snmp mib Output Fields	361
	Table 23: show snmp rmon Output Fields	363
	Table 24: show smp rmon history Output Fields	367
	Table 25: show snmp statistics Output Fields	371
	Table 26: show snmp v3 Output Fields	376
Chapter 19	Operational Commands: Port Mirroring	379
	Table 27: show analyzer Output Fields	380
Chapter 20	Operational Commands: sFlow	383
	Table 28: show sflow Output Fields	384
	Table 29: show sflow interface Output Fields	386
	Table 30: show sflow collector Output Fields	388
Chapter 21	Operational Commands: Ethernet OAM Connectivity Fault Management	391
	Table 31: monitor ethernet delay-measurement one-way Output Fields	398
	Table 32: monitor ethernet delay-measurement two-way Output Fields	398
	Table 33: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields	402
	Table 34: show oam ethernet connectivity-fault-management forwarding-state Output Fields	405
	Table 35: show oam ethernet connectivity-fault-management interfaces Output Fields	409
	Table 36: show oam ethernet connectivity-fault-management linktrace path-database Output Fields	415
	Table 37: show oam ethernet connectivity-fault-management mep-database Output Fields	418
	Table 38: show oam ethernet connectivity-fault-management mip Output Fields	424
	Table 39: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields	425
Chapter 22	Operational Commands: Ethernet OAM Link Fault Management	431
	Table 40: show oam ethernet link-fault-management Output Fields	432
Chapter 23	Operational Commands: Uplink Failure Detection	437
	Table 41: show uplink-failure-detection Output Fields	438
Part 4	Troubleshooting	
Chapter 24	Troubleshooting Procedures and Commands	443
	Table 42: show pfe statistics bridge Output Fields	446

About the Documentation

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

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:

- EX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xix 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 xix 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 at the Juniper Networks Technical Documentation 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 <https://www.juniper.net/cgi-bin/docbugreport/>.
- 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 JNASC 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

- [Real-Time Performance Monitoring on page 3](#)
- [Port Mirroring on page 9](#)
- [sFlow Monitoring Technology on page 17](#)
- [Ethernet OAM Link Fault Management and Connectivity Fault Management on page 21](#)
- [Uplink Failure Detection on page 27](#)

CHAPTER 1

Real-Time Performance Monitoring

- [Understanding Real-Time Performance Monitoring on EX Series Switches on page 3](#)

Understanding Real-Time Performance Monitoring on EX Series Switches

Real-time performance monitoring (RPM) enables you to configure active probes to track and monitor traffic across the network and to investigate network problems. You can use RPM with Juniper Networks EX Series Ethernet Switches.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.
- Set thresholds to trigger SNMP traps when values are exceeded.

You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test. (SNMP trap results are stored in `pingResultsTable`, `jnxPingResultsTable`, `jnxPingProbeHistoryTable`, and `pingProbeHistoryTable`.)

- Determine automatically whether a path exists between a host router or switch and its configured BGP neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

RPM provides MIB support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

This topic includes:

- [RPM Packet Collection on page 4](#)
- [Tests and Probe Types on page 4](#)
- [Hardware Timestamps on page 4](#)
- [Limitations of RPM on EX Series Switches on page 6](#)

RPM Packet Collection

Probes collect packets per destination and per application, including ping Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets.

Tests and Probe Types

A test can contain multiple probes. The probe type specifies the packet and protocol contents of the probe.

EX Series switches support the following tests and probe types:

- Ping tests:
 - ICMP echo probe
 - ICMP timestamp probe
- HTTP tests:
 - HTTP get probe (not available for BGP RPM services)
 - HTTP get metadata probe
- UDP and TCP tests with user-configured ports:
 - UDP echo probe
 - TCP connection probe
 - UDP timestamp probe

Hardware Timestamps

To account for latency or jitter in the communication of probe messages, you can enable timestamping of the probe packets (hardware timestamps). If hardware timestamps are not configured, then timers are generated at the software level and are less accurate than they would have been with hardware timestamps.



NOTE: EX Series switches support hardware timestamps for UDP and ICMP probes. EX Series switches do not support hardware timestamps for HTTP or TCP 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


You should configure the requester (the RPM client) with hardware timestamps (see [Figure 1 on page 5](#)) to get more meaningful results than you would get without the timestamps. The responder (the RPM server) does not need to be configured to support hardware timestamps. If the responder supports hardware timestamps, it timestamps the RPM probes. If the responder does not support hardware timestamps, RPM can only report round-trip measurements that include the processing time on the responder.



NOTE: Hardware timestamps are supported on all EX Series switches.

[Figure 1 on page 5](#) shows the timestamps:

Figure 1: RPM Timestamps



- T1 is the time the packet leaves the requester port.
- T2 is the time the responder receives the packet.
- T3 is the time the responder sends the response.
- T4 is the time the requester receives the response.

The round-trip time is $(T2 - T1) + (T4 - T3)$. If the responder does not support hardware timestamps, then the round-trip time is $(T4 - T1) / 2$, and thus includes the processing time of the responder.

You can use RPM probes to find the following time measurements:

- Minimum round-trip time
- Maximum round-trip time
- Average round-trip time
- Standard deviation of the round-trip time
- Jitter of the round-trip time—Difference between the minimum and maximum round-trip time



NOTE: See “[Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch \(CLI Procedure\)](#)” on [page 92](#) for information on how to configure hardware timestamps on the requester.

The RPM feature provides a configuration option to set one-way hardware timestamps. Use one-way timestamps when you want information about one-way time, rather than round-trip times, for packets to traverse the network between the requester and the responder. As shown in [Figure 1 on page 5](#), one-way timestamps represent the time $T2 - T1$ and the time from $T4 - T3$. Use one-way timestamps when you want to gather information about delay in each direction and to find egress and ingress jitter values.



NOTE: For correct one-way measurement, the clocks of the requester and responder must be synchronized. If the clocks are not synchronized, one-way jitter measurements and calculations can include significant variations, in some cases orders of magnitude greater than the round-trip times.

When you enable one-way timestamps in a probe, the following one-way measurements are reported:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probes sent
- Number of probe responses received
- Percentage of lost probes

Limitations of RPM on EX Series Switches

- Two-Way Active Measurement Protocol (TWAMP) is not supported on EX Series switches.
- EX Series switches do not support user-configured class-of-service (CoS) classifiers or prioritization of RPM packets over regular data packets received on an input interface.
- Timestamps:
 - If the responder does not support hardware timestamps, RPM can only report the round-trip measurements and cannot calculate round-trip jitter.
 - EX Series switches do not support hardware timestamps for HTTP and TCP probes.
 - Timestamps apply only to IPv4 traffic.

Related Documentation

- [For further details about RPM, see *Junos OS Services Interfaces Configuration Guide*](#)
- [Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch \(CLI Procedure\) on page 92](#)
- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 85](#)

- [Configuring SNMP \(J-Web Procedure\) on page 75](#)
- [Monitoring Network Traffic Using Traceroute on page 299](#)

CHAPTER 2

Port Mirroring

- [Understanding Port Mirroring on EX Series Switches on page 9](#)

Understanding Port Mirroring on EX Series Switches

You can use port mirroring to facilitate analyzing traffic on your Juniper Networks EX Series Ethernet Switch on a packet level. You might use port mirroring as part of monitoring switch traffic for such purposes as enforcing policies concerning network usage and file sharing and for identifying sources of problems on your network by locating abnormal or heavy bandwidth usage by particular stations or applications.

Port mirroring copies packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on Juniper Networks EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, or EX6200 Ethernet Switches
- Packets exiting a VLAN on Juniper Networks EX8200 Ethernet Switches

This topic describes:

- [Port Mirroring Overview on page 9](#)
- [Port Mirroring Terminology on page 10](#)
- [Configuration Guidelines for Port Mirroring on the Switches on page 12](#)

Port Mirroring Overview

Port mirroring might be needed for traffic analysis on a switch because a switch, unlike a hub, does not broadcast packets to every port on the destination device. The switch sends packets only to the port to which the destination device is connected.

You configure port mirroring on the switch to send copies of unicast traffic to either a local analyzer port or an analyzer VLAN. Then you can analyze the mirrored traffic using a protocol analyzer application. The protocol analyzer application can run either on a computer connected to the analyzer output interface or on a remote monitoring station.

You can use port mirroring on a switch to mirror any of the following:

- **Packets entering or exiting a port**—You can mirror the packets in any combination (on up to 256 ports). For example, you can send copies of the packets entering some ports and the packets exiting other ports to the same local analyzer port or analyzer VLAN.
- **Packets entering a VLAN on an EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, or EX6200 switch**—You can mirror the packets entering a VLAN on these switches to either a local analyzer port or to an analyzer VLAN. On EX3200, EX4200, EX4500, and EX4550 switches, you can configure multiple VLANs (up to 256 VLANs), including a VLAN range and PVLANS, as ingress input to an analyzer.
- **Packets exiting a VLAN on an EX8200 switch**—You can mirror the packets exiting a VLAN on an EX8200 switch to either a local analyzer port or to an analyzer VLAN. You can configure multiple VLANs (up to 256 VLANs), including a VLAN range and PVLANS, as egress input to an analyzer.
- **Statistical samples**—You can mirror a statistical sample of packets that are
 - Entering or exiting a port
 - Entering a VLAN on an EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, or EX6200 switch
 - Exiting a VLAN on an EX8200 switch

You specify the sample number of packets by setting the ratio. You can send the sample to either a local analyzer port or to an analyzer VLAN.

- **Policy-based sample**—You can mirror a policy-based sample of packets that are entering a port or a VLAN. You configure a firewall filter to establish a policy to select the packets to be mirrored. You can send the sample to a local analyzer port or to an analyzer VLAN.



NOTE: Juniper Networks Junos operating system (Junos OS) for EX Series switches implements port mirroring differently than other Junos OS packages. Junos OS for EX Series switches does not include the `port-mirroring` statement found in the `edit forwarding-options` level of the hierarchy of other Junos OS packages, or the `port-mirror` action in firewall filter terms.

Port Mirroring Terminology

Table 3 on page 11 lists some port mirroring terms and their descriptions.

Table 3: Port Mirroring Terminology

Term	Description
Analyzer	<p>A port mirroring configuration on an EX Series switch. The analyzer includes:</p> <ul style="list-style-type: none"> • The name of the analyzer • Source (input) ports or VLAN (optional) • A destination for mirrored packets (either a monitor port or a monitor VLAN) • Ratio field for specifying statistical sampling of packets (optional) • Loss-priority setting
Analyzer output interface (Also known as monitor port)	<p>Interface to which mirrored traffic is sent and to which a protocol analyzer application is connected.</p> <p>NOTE: Interfaces used as output for an analyzer must be configured as family ethernet-switching.</p> <p>Analyzer output interfaces have the following limitations:</p> <ul style="list-style-type: none"> • Cannot also be a source port. • Cannot be used for switching. • Do not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP), when part of a port mirroring configuration. • Do not retain any VLAN associations they held before they were configured as analyzer output interfaces. <p>If the bandwidth of the analyzer output interface is not sufficient to handle the traffic from the source ports, overflow packets are dropped.</p>
Analyzer VLAN (Also known as monitor VLAN)	<p>VLAN to which mirrored traffic is sent. The mirrored traffic can be used by a protocol analyzer application. The member interfaces in the monitor VLAN are spread across the switches in your network.</p>
Firewall-based analyzer	<p>An analyzer whose configuration does not specify an input source; it specifies only an output destination. A firewall-based analyzer must be used with a firewall filter to achieve the functionality of an analyzer.</p>
Global analyzer (on EX4500 and EX4550 switches only)	<p>An analyzer that is based on a firewall filter, VLAN, or link aggregation group (LAG) or an analyzer in which interfaces are on different port groups on the switch. A port group is a logical group of ports on the switch.</p>
Input interface (Also known as mirrored ports or monitored interfaces)	<p>An interface on the switch that is being mirrored, on traffic that is either entering or exiting the interface. An input interface cannot also be an output interface for an analyzer.</p>
LAG-based analyzer	<p>An analyzer that has a LAG specified as the input (ingress) interface in the analyzer configuration.</p>
Local port mirroring	<p>An analyzer configuration in which packets are mirrored to a local analyzer port.</p>
Mirror ratio	<p>See statistical sampling.</p>
Monitoring station	<p>A computer running a protocol analyzer application.</p>

Table 3: Port Mirroring Terminology (*continued*)

Term	Description
Native analyzer session	An analyzer session that has both input and output definitions in its analyzer configuration.
Policy-based mirroring	Mirroring of packets that match the match items in the defined firewall filter term. The action item analyzer analyzer-name is used in the firewall filter to send the packets to the analyzer.
Port-based analyzer	An analyzer session whose configuration defines interfaces for both input and output.
Protocol analyzer application	An application used to examine packets transmitted across a network segment. Also commonly called network analyzer, packet sniffer, or probe.
Remote port mirroring	<p>Functions the same as local port mirroring, except that the mirrored traffic is not copied to a local analyzer port but is flooded to an analyzer VLAN that you create specifically for the purpose of receiving mirrored traffic.</p> <p>If you are using an intermediate (transit) switch, you can avoid flooding of the mirrored traffic to member interfaces of the VLAN by setting the ingress option to specify an interface of the VLAN for ingress-only traffic and the egress option to specify an interface of the VLAN for egress-only traffic in the [edit vlans] hierarchy level. See "Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches" on page 61.</p>
Statistical sampling	<p>You can configure the system to mirror a sampling of the packets by setting a ratio of 1:x, where x is a value from 1 through 2047.</p> <p>For example, when x is set to 1, all packets are copied to the analyzer. When x is set to 200, 1 of every 200 packets is copied.</p>
VLAN-based analyzer	An analyzer session whose configuration uses VLANs for both input and output or for either input or output.

Configuration Guidelines for Port Mirroring on the Switches

When you configure port mirroring, we recommend that you follow certain guidelines to ensure that you obtain optimum benefit from the port mirroring feature. Additionally, we recommend that you disable port mirroring when you are not using it and that you select specific interfaces for which packets must be mirrored (that is, select specific interfaces as input to the analyzer) in preference to using the **all** keyword option, which will enable port mirroring on all interfaces. You can also limit the amount of mirrored traffic by using statistical sampling, setting a ratio to select a statistical sample, or using a firewall filter. Mirroring only the necessary packets reduces any potential performance impact.

With local port mirroring, traffic from multiple ports is replicated to the analyzer output interface. If the output interface for an analyzer reaches capacity, packets are dropped. Thus, while configuring an analyzer, you must consider whether the traffic being mirrored exceeds the capacity of the analyzer output interface.

[Table 4 on page 13](#) summarizes further configuration guidelines for port mirroring on the switches.

Table 4: Configuration Guidelines for Port Mirroring

Guideline	Description	Comment
<p>NOTE: “All other switches” or “All switches” in the Description column applies to switch platforms that support port mirroring. For details on platform support, see <i>EX Series Switch Software Features Overview</i>.</p>		
Number of VLANs that you can use as ingress input to an analyzer	<ul style="list-style-type: none"> 1—EX2200 switches 256—EX3200, EX4200, EX4500, EX4550, and EX6200 switches Does not apply—EX8200 switches 	
Number of analyzers that you can enable concurrently	<ul style="list-style-type: none"> 1—EX2200, EX3200, EX4200, EX3300, and EX6200 switches 7 port-based or 1 global—EX4500 and EX4550 switches 7 total, with one based on a VLAN, firewall filter, or LAG and with the remaining 6 based on firewall filters—EX8200 switches <p>NOTE: An analyzer configured using a firewall filter does not support mirroring of packets that are egressing ports.</p>	<ul style="list-style-type: none"> You can <i>configure</i> more than the specified number of analyzers on the switch, but you can <i>enable</i> only the specified number for a session. Use disable ethernet-switching-options analyzer name to disable an analyzer. See Table 3 on page 11 for a description of global analyzers. See the next row entry in this table for the exception to the number of firewall-filter-based analyzers allowed on EX4500 and EX4550 switches. On an EX4550 Virtual Chassis, you can configure only one analyzer if ports in the input and output definitions are on different switches in a Virtual Chassis. To configure multiple analyzers, an entire analyzer session must be configured on the same switch of a Virtual Chassis.
Number of firewall-filter-based analyzers that you can configure on EX4500 and EX4550 switches	<ul style="list-style-type: none"> 1—EX4500 and EX4550 switches 	If you configure multiple analyzers, you cannot attach any of them to a firewall filter.
Types of ports on which you cannot mirror traffic	<ul style="list-style-type: none"> Virtual Chassis ports (VCPs) Management Ethernet ports (me0 or vme0) Routed VLAN interfaces (RVIs) VLAN-tagged Layer 3 interfaces 	
If port mirroring is configured to mirror packets exiting 10-Gigabit Ethernet ports on EX8200 switches, packets are dropped in both network and mirrored traffic when the mirrored packets exceed 60 percent of the 10-Gigabit Ethernet port traffic.	<ul style="list-style-type: none"> EX8200 switches 	
Traffic directions for which you can specify a ratio	<ul style="list-style-type: none"> Ingress only—EX8200 switches Ingress and egress—All other switches 	

Table 4: Configuration Guidelines for Port Mirroring (*continued*)

Guideline	Description	Comment
Protocol families that you can include in a firewall-filter-based remote analyzer	<ul style="list-style-type: none"> Any except inet and inet6—EX8200 switches Any—All other switches 	You can use inet and inet6 on EX8200 switches in a local analyzer.
Traffic directions that you can configure for mirroring on ports in firewall-filter-based configurations	<ul style="list-style-type: none"> Ingress only—All switches 	
Mirrored packets on tagged interfaces might contain an incorrect VLAN ID or Ethertype.	<ul style="list-style-type: none"> Both VLAN ID and Ethertype—EX2200 switches VLAN ID only—EX3200 and EX4200 switches Ethertype only—EX4500 and EX4550 switches Does not apply—EX8200 switches 	
Mirrored packets exiting an interface do not reflect rewritten class-of-service (CoS) DSCP or 802.1p bits.	<ul style="list-style-type: none"> All switches 	
The analyzer appends an incorrect 802.1Q (dot1q) header to the mirrored packets on the routed traffic or does not mirror any packets on the routed traffic when an egress VLAN that belongs to a routed VLAN interface (RVI) is configured as the input for that analyzer.	<ul style="list-style-type: none"> EX8200 switches Does not apply—All other switches 	As a workaround, configure an analyzer that uses each port (member interface) of the VLAN as egress input.
Packets with physical layer errors are not sent to the local or remote analyzer.	<ul style="list-style-type: none"> All switches 	Packets with these errors are filtered out and thus are not sent to the analyzer.
Port mirroring configuration on a Layer 3 interface with the output configured to a VLAN is not available on EX8200 switches.	<ul style="list-style-type: none"> EX8200 switches Does not apply—All other switches 	
Port mirroring does not support line-rate traffic.	<ul style="list-style-type: none"> All switches 	Port mirroring for line-rate traffic is done on a best-effort basis.
In an EX8200 Virtual Chassis, if you need to mirror traffic across the virtual chassis, then the output port must be a LAG.	<ul style="list-style-type: none"> EX8200 Virtual Chassis Does not apply—All other switches 	<p>In an EX8200 Virtual Chassis:</p> <ul style="list-style-type: none"> You can configure LAG as a monitor port only for native analyzers. You cannot configure LAG as a monitor port for analyzers based on firewall filters. If an analyzer configuration contains LAG as a monitor port, then you cannot configure VLAN in the input definition of an analyzer.

Table 4: Configuration Guidelines for Port Mirroring (*continued*)

Guideline	Description	Comment
In standalone EX8200 switches, you can configure LAG in the output definition.	<ul style="list-style-type: none">EX8200 standalone switchesDoes not apply—All other switches	<p>In EX8200 standalone switches:</p> <ul style="list-style-type: none">You can configure a LAG as a monitor port on both native and firewall-based analyzers.If a configuration contains LAG as a monitor port, then you cannot configure VLAN in the input definition of an analyzer.

- Related Documentation
- [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37](#)
 - [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
 - [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
 - [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)
 - *Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches*

CHAPTER 3

sFlow Monitoring Technology

- [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 17](#)

Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. You can configure sFlow technology on a Juniper Networks EX Series Ethernet Switch to continuously monitor traffic at wire speed on all interfaces simultaneously.

This topic describes:

- [Sampling Mechanism and Architecture of sFlow Technology on EX Series Switches on page 17](#)
- [Adaptive Sampling on page 18](#)
- [sFlow Agent Address Assignment on page 19](#)

Sampling Mechanism and Architecture of sFlow Technology on EX Series Switches

sFlow technology uses the following two sampling mechanisms:

- **Packet-based sampling:** Samples one packet out of a specified number of packets from an interface enabled for sFlow technology.
- **Time-based sampling:** Samples interface statistics at a specified interval from an interface enabled for sFlow technology.

The sampling information is used to create a network traffic visibility picture. The Juniper Networks Junos operating system (Junos OS) fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.



NOTE: sFlow technology on the switches samples only raw packet headers. A raw Ethernet packet is the complete Layer 2 network frame.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent's two main activities are random sampling and statistics gathering. The sFlow agent combines interface counters and flow samples and sends them across the network to the sFlow collector in UDP datagrams, directing those datagrams to the IP address and UDP destination port of the collector. Each datagram contains the following information:

- The IP address of the sFlow agent
- The number of samples
- The interface through which the packets entered the agent
- The interface through which the packets exited the agent
- The source and destination interface for the packets
- The source and destination VLAN for the packets

EX Series switches adopt the distributed sFlow architecture. The sFlow agent has two separate sampling entities that are associated with each Packet Forwarding Engine. These sampling entities are known as subagents. Each subagent has a unique ID that is used by the collector to identify the data source. A subagent has its own independent state and forwards its own sample packets to the sFlow agent. The sFlow agent is responsible for packaging the samples into datagrams and sending them to the sFlow collector. Because sampling is distributed across subagents, the protocol overhead associated with sFlow technology is significantly reduced at the collector.



NOTE: You cannot configure sFlow monitoring on a link aggregation group (LAG), but you can configure it individually on a LAG member interface.



NOTE: If the mastership assignment changes in a Virtual Chassis setup, sFlow technology continues to function.

Adaptive Sampling

The switches use adaptive sampling to ensure both sampling accuracy and efficiency. Adaptive sampling is a process of monitoring the overall incoming traffic rate on the network device and providing intelligent feedback to interfaces to dynamically adapt the sampling rates on interfaces on the basis of traffic conditions. Interfaces on which incoming traffic exceeds the system threshold are checked so that all violations can be regulated without affecting the traffic on other interfaces. Every 12 seconds, the agent checks interfaces to get the number of samples, and interfaces are grouped on the basis of the slot that they belong to. The top five interfaces that produce the highest number of samples are selected. Using the binary backoff algorithm, the sampling load on these interfaces is reduced by half and allotted to interfaces that have a lower sampling rate. Therefore, when the processor's sampling limit is reached, the sampling rate is adapted such that it does not load the processor any further. If the switch is rebooted, the adaptive

sampling rate is reset to the user-configured sampling rate. Also, if you modify the sampling rate, the adaptive sampling rate changes.

The advantage of adaptive sampling is that the switch continues to operate at its optimum level even when there is a change in the traffic patterns in the interfaces. You do not need to make any changes. Because the sampling rate adapts dynamically to changing network conditions, the resources are utilized optimally resulting in a high-performance network.

Infrequent sampling flows might not be reported in the sFlow information, but over time, the majority of flows are reported. On the basis of the configured sampling rate N , 1 out of N packets is captured and sent to the collector. This type of sampling does not provide a result that is 100 percent accurate in the analysis, but it does provide a result of quantifiable accuracy. A user-configured polling interval defines how often the sFlow data for a specific interface are sent to the collector, but an sFlow agent can also schedule polling.



NOTE: sFlow technology on EX Series switches does not support graceful restart. When a graceful restart occurs, the adaptive sampling rate is set to the user-configured sampling rate.

sFlow Agent Address Assignment

The sFlow collector uses the sFlow agent's IP address to determine the source of the sFlow data. You can configure the IP address of the sFlow agent to ensure that the agent ID of the sFlow agent remains constant. If you do not configure the IP address of the sFlow agent, an IP address is automatically assigned to the agent. This is the IP address of one of the following interfaces configured on the switch taken in the given order of priority:

1. Virtual management Ethernet (VME) interface
2. Management Ethernet interface

If neither of the preceding interfaces has been configured, the IP address of any Layer 3 interface or the routed VLAN interface (RVI) is assigned to the agent. At least one interface must be configured on the switch for an IP address to be automatically assigned to the agent. When the agent's IP address is assigned automatically, the IP address is dynamic and changes when the switch reboots.

sFlow data can be used to provide network traffic visibility information. You can explicitly configure the IP address to be assigned to source data (sFlow datagrams). If you do not explicitly configure that address, the IP address of the configured Gigabit Ethernet interface, 10-Gigabit Ethernet interface, or the RVI is used as the source IP address.

Related Documentation

- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 84](#)
- [Monitoring Interface Status and Traffic](#)

CHAPTER 4

Ethernet OAM Link Fault Management and Connectivity Fault Management

- [Understanding Ethernet OAM Link Fault Management for an EX Series Switch on page 21](#)
- [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch on page 22](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 24](#)

Understanding Ethernet OAM Link Fault Management for an EX Series Switch

Juniper Networks Junos operating system (Junos OS) for Juniper Networks EX Series Ethernet Switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported on EX Series switches:

- Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer

on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote Fault Detection

Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection. The EX Series switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote Loopback Mode

Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

**Related
Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 97](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72](#)

Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch

Ethernet interfaces on Juniper Networks EX Series Ethernet Switches and Juniper Networks Junos operating system (Junos OS) for EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). CFM monitors Ethernet networks that might comprise one or more service instances for network-compromising connectivity faults.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol.
- Fault isolation using the loopback protocol.


CFM partitions the service network into various administrative domains. For example, operators, providers, and customers might be part of different administrative domains. Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible.

In a CFM maintenance domain, each service instance is called a maintenance association. A maintenance association can be thought of as a full mesh of maintenance association endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages. There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outer domains are assigned a higher level than the inner domains. Configure customer end points to have the highest maintenance domain level. The maintenance domain level is a mandatory parameter that indicates the nesting relationships between various maintenance domains. The level is embedded in each CFM frame. CFM messages within a given level are processed by MEPs at that same level.

To enable CFM on an Ethernet interface, you must configure maintenance domains, maintenance associations, and maintenance association end points (MEPs). [Figure 2 on page 24](#) shows the relationships among maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs) configured on a switch.

Figure 2: Relationship Among MEPs, MIPs, and Maintenance Domain Levels



Related Documentation

- [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 93](#)
- [Junos OS Network Interfaces Configuration Guide](#)

Understanding Ethernet Frame Delay Measurements on Switches

Performance management depends on the accurate measurement of service-level agreement (SLA) objective parameters, which can include bandwidth and reliability. In many cases, a service provider could be subject to penalties imposed by regulation, statute, or contract if network performance is not within the bounds established for the service. One key performance objective is delay, along with its close relative, delay variation (often called jitter). Some applications (such as bulk file transfer) will function just as well with high delays across the network and high delay variations, while other applications (such as voice) can function only with low and stable delays. Many networks invoke protocols or features available at Layer 3 (the packet layer) or higher to measure network delays and jitter link by link. However, when the network consists of many Ethernet links, there are few protocols and features available at Layer 2 (the frame layer) that allow routers and switches to measure frame delay and jitter. This is where the ability to configure and monitor Ethernet frame delay is helpful.

This topic includes:

- [Ethernet Frame Delay Measurements on page 24](#)
- [Types of Ethernet Frame Delay Measurements on page 25](#)
- [Limitations on page 26](#)

Ethernet Frame Delay Measurements

You can perform Ethernet frame delay measurements (referred to as ETH-DM in Ethernet specifications) on Juniper Networks EX Series Ethernet Switches. This feature allows you to configure on-demand Operation, Administration, and Maintenance (OAM)

statements for the measurement of frame delay and frame delay variation (jitter). You can configure Ethernet frame delay measurement in either one-way or two-way (round-trip) mode to gather frame delay statistics simultaneously from multiple sessions. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor SLAs.

Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. It supports software-assisted timestamping in the receive direction for delay measurements. It also provides runtime display of delay statistics when two-way delay measurement is triggered. Ethernet frame delay measurement records the last 100 samples collected per remote maintenance association end point (MEP) or per connectivity fault management (CFM) session. You can retrieve the history at any time using simple commands. You can clear all Ethernet frame delay measurement statistics and PDU counters. Ethernet frame delay measurement is fully compliant with the ITU-T Y.1731 (*OAM Functions and Mechanisms for Ethernet-based Networks*) specification.

Ethernet frame delay measurement uses the IEEE 802.1ag CFM infrastructure.

Generally, Ethernet frame delay measurements are made in a peer fashion from one MEP or CFM session to another. However, these measurements are not made to maintenance association intermediate points (MIPs).

For a complete description of Ethernet frame delay measurement, see the *ITU-T Y.1731 Ethernet Service OAM* topics in the *Junos OS Network Interfaces Library for Routing Devices*.

Types of Ethernet Frame Delay Measurements

There are two types of Ethernet frame delay measurements:

- One-way
- Two-way (round-trip)

For one-way Ethernet frame delay measurement, either MEP can send a request to begin a one-way delay measurement to its peer MEP. However, the statistics are collected only at the receiver MEP. This feature requires the clocks at the transmitting and receiving MEPs to be synchronized. If these clocks fall out of synchronization, only one-way delay variation and average delay variation values are computed correctly (and will, therefore, be valid). Use the **show** commands at the receiver MEP to display one-way delay statistics.

For two-way (round-trip) Ethernet frame delay measurement, either MEP can send a request to begin a two-way delay measurement to its peer MEP, which responds with timestamp information. Run-time statistics are collected and displayed at the initiator MEP. The clocks do not need to be synchronized at the transmitting and receiving MEPs. Junos OS supports timestamps in delay measurement reply (DMR) frames to increase the accuracy of delay calculations.

Use the **show** commands at the initiator MEP to display two-way delay statistics, and at the receiver MEP to display one-way delay statistics.

You can create an iterator profile to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for delay measurement or loss measurement.

Limitations

The following are some limitations with regard to using Ethernet frame delay measurement:

- Ethernet frame delay measurements are available only when distributed periodic packet management (PPM) is enabled.
- The statistics collected are lost after a graceful Routing Engine switchover (GRES).
- You can monitor only one session to the same remote MEP or MAC address.
- Accuracy is compromised when the system configuration changes (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 103](#)

CHAPTER 5

Uplink Failure Detection

- [Understanding Uplink Failure Detection on page 27](#)

Understanding Uplink Failure Detection

Uplink failure detection allows Juniper Networks EX Series Ethernet Switches to detect link failure on uplink interfaces and to propagate the failure to the downlink interfaces so that servers connected to those downlink interfaces can switch over to secondary interfaces.

Uplink failure detection supports network adapter teaming and provides network redundancy. In network adapter teaming, all the network interface cards (NICs) on a server are configured in a primary or secondary relationship and share the same IP address. When the primary link goes down, the server transparently shifts the connection to the secondary link. With uplink failure detection, the switch monitors uplink interfaces for link failures. When it detects a failure, it disables the downlink interfaces. When the server detects disabled downlink interfaces, it switches over to the secondary link to help ensure balanced traffic flow on switches.

This topic describes:


- [Uplink Failure Detection Overview on page 27](#)
- [Failure Detection Pair on page 28](#)

Uplink Failure Detection Overview

Uplink failure detection allows switches to monitor uplink interfaces to spot link failures. When a switch detects a link failure, it automatically disables the downlink interfaces in that group. The server that is connected to the disabled downlink interfaces triggers a network-adapter failover to a secondary link to avoid any information drop.

[Figure 3 on page 28](#) illustrates a typical setup for uplink failure detection.

Figure 3: Uplink Failure Detection Configuration on Switches



For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down all downlink interfaces within that group. If any uplink interface returns to service, then the switch brings all downlink interfaces in that group back to service.



NOTE: Routed VLAN interfaces (RVIs) cannot be configured as uplink interfaces to be monitored.

The switch can monitor both physical-interface links and logical-interface links for uplink failures, but you must put the two types of interfaces in separate groups.



NOTE: To detect failure of logical interfaces, the server must run some high level protocol such as keepalives between the switch and the server.

Failure Detection Pair

Uplink failure detection requires that you create groups that contain uplink interfaces and downlink interfaces. Each group includes one of each of the following:

- A link-to-monitor interface—The link-to-monitor interfaces specify the uplink interfaces the switch monitors. You can configure a maximum of 48 uplink interfaces as link-to-monitor in a group.

- A link-to-disable interface—The link-to-disable interfaces specify the downlink interfaces the switch disables when the switch detects an uplink failure. You can configure a maximum of 48 downlink interfaces as link-to-disable in a group.

The link-to-disable interfaces are bound to the link-to-monitor interfaces within the group. When a link-to-monitor interface returns to service, the switch automatically enables all link-to-disable interfaces in the group.

**Related
Documentation**

- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 99](#)

PART 2

Configuration

- [Configuration Examples on page 33](#)
- [Configuration Tasks on page 75](#)
- [Configuration Statements: SNMP on page 105](#)
- [Configuration Statements: Port Mirroring on page 185](#)
- [Configuration Statements: sFlow Technology on page 203](#)
- [Configuration Statements: Ethernet OAM Connectivity Fault Management on page 213](#)
- [Configuration Statements: Ethernet OAM Link Fault Management on page 245](#)
- [Configuration Statements: RPM on page 267](#)
- [Configuration Statements: Uplink Failure Detection on page 289](#)

CHAPTER 6

Configuration Examples

- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33](#)
- [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37](#)
- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
- [Example: Configuring Port Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX Series Switches on page 51](#)
- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61](#)
- [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 68](#)
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72](#)

Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches

sFlow technology is a networking monitoring technology for high-speed switched or routed networks. It is a technology that is based on statistical sampling. You can configure sFlow technology to continuously monitor traffic at wire speed on all interfaces simultaneously. sFlow data can be used to provide network traffic visibility information. You can specify sampling rates for ingress and egress packets. Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

This example describes how to configure and use sFlow technology to monitor network traffic.

- [Requirements on page 34](#)
- [Overview and Topology on page 34](#)
- [Configuration on page 34](#)
- [Verification on page 36](#)

Requirements

This example uses the following hardware and software components:

- One EX Series switch
- Junos OS Release 9.3 or later for EX Series switches


Overview and Topology

sFlow technology samples network packets and sends the samples to a monitoring station. You can specify sampling rates for ingress and egress packets. The information gathered is used to create a network traffic visibility picture.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent runs on the switch. It combines interface counters and flow samples and sends them across the network to the sFlow collector.

[Figure 4 on page 34](#) depicts the basic elements of the sFlow system.

Figure 4: sFlow Technology Monitoring System



Configuration

To configure sFlow technology, perform the following tasks:

CLI Quick Configuration To quickly configure sFlow technology, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]
set sflow collector 10.204.32.46 udp-port 5600
set sflow interfaces ge-0/0/0
set sflow polling-interval 20
set sflow sample-rate egress 1000
```

Step-by-Step Procedure To configure sFlow technology:

1. Configure the IP address and UDP port of the collector:

```
[edit protocols]
user@switch# set sflow collector 10.204.32.46 udp-port 5600
```



NOTE: You can configure a maximum of 4 collectors.

The default UDP port is 6343.

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces ge-0/0/0
```



NOTE: You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a link aggregation group (LAG) interface, but you can enable it on the member interfaces of a LAG.

3. Specify in seconds how often the sFlow agent polls the interface:

```
[edit protocols sflow]
user@switch# set polling-interval 20
```



NOTE: The polling interval can be specified as a global parameter also. Specify 0 if you do not want to poll the interface.

4. Specify the rate at which egress packets must be sampled:

```
[edit protocols sflow]
user@switch# set sample-rate egress 1000
```



NOTE: You can specify both egress and ingress sampling rates. If you set only the egress sampling rate, the ingress sampling rate will be disabled.



NOTE: We recommend that you configure the same sampling rates on all the ports on a line card. If you configure different sampling rates are different, the lowest value is used for all ports. You could still configure different rates on different line cards.

Results Check the results of the configuration:

```
[edit protocols sflow]
user@switch# show
polling-interval 20;
sample-rate egress 1000;
collector 10.204.32.46 {
  udp-port 5600;
}
interfaces ge-0/0/0.0;
```

Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That sFlow Technology Is Configured Properly on page 36](#)
- [Verifying That sFlow Technology Is Enabled on the Specified Interface on page 37](#)
- [Verifying the sFlow Collector Configuration on page 37](#)

Verifying That sFlow Technology Is Configured Properly

Purpose Verify that sFlow technology is configured properly.

Action Use the **show sflow** command:

```
user@switch> show sflow
sFlow: Enabled
Sample limit: 300 packets/second
Polling interval: 20 seconds
Sample rate egress: 1:1000: Enabled
Sample rate ingress: 1:2048: Disabled
Agent ID: 10.204.96.222
```



NOTE: The sampling limit cannot be configured and is set to 300 packets/second per FPC.

Meaning The output shows that sFlow technology is enabled and specifies the values for the sampling limit, polling interval, and the egress sampling rate.

Verifying That sFlow Technology Is Enabled on the Specified Interface

- Purpose** Verify that sFlow technology is enabled on the specified interfaces and display the sampling parameters.
- Action** Use the `show sflow interface` command:
- ```
user@switch> show sflow interface
```
- | Interface  | Status           | Sample rate    | Adapted sample rate | Polling-interval |
|------------|------------------|----------------|---------------------|------------------|
|            | Egress Ingress   | Egress Ingress | Egress Ingress      |                  |
| ge-0/0/0.0 | Enabled Disabled | 1000 2048      | 1000 2048           | 20               |
- Meaning** The output indicates that sFlow technology is enabled on the ge-0/0/0.0 interface with an egress sampling rate of 1000, a disabled ingress sampling rate, and a polling interval of 20 seconds.

### Verifying the sFlow Collector Configuration

- Purpose** Verify the sFlow collector's configuration.
- Action** Use the `show sflow collector` command:
- ```
user@switch> show sflow collector
```
- | Collector address | Udp-port | No. of samples |
|-------------------|----------|----------------|
| 10.204.32.46 | 5600 | 1000 |
| 10.204.32.76 | 3400 | 1000 |
- Meaning** The output displays the IP address of the collectors and the UDP ports. It also displays the number of samples.
- Related Documentation**
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 84](#)
 - [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 17](#)

Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, or EX6200 switches
- Packets exiting a VLAN on EX8200 switches

You can analyze the mirrored traffic using a protocol analyzer application installed on a system connected to the local destination interface (or a running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN).

This example describes how to configure an EX Series switch to mirror traffic entering interfaces connected to employee computers to an analyzer output interface on the same switch.

This example describes how to configure local port mirroring:

- [Requirements on page 38](#)
- [Overview and Topology on page 38](#)
- [Mirroring All Employee Traffic for Local Analysis on page 39](#)
- [Mirroring Employee-to-Web Traffic for Local Analysis on page 40](#)
- [Verification on page 42](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.0 or later for EX Series switches
- One EX Series switch

Before you configure port mirroring, be sure you have an understanding of port mirroring concepts.

Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to a destination interface on the same switch. The first example shows how to mirror all traffic entering the ports connected to employee computers. The second example shows the same scenario, but includes a filter to mirror only the employee traffic going to the Web.

Network Topology

In this example, **ge-0/0/0** and **ge-0/0/1** serve as connections for employee computers.


In this example, one interface, **ge-0/0/10**, is reserved for analysis of mirrored traffic. Connect a PC running a protocol analyzer application to the analyzer output interface to analyze the mirrored traffic.



NOTE: Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

[Figure 5 on page 39](#) shows the network topology for this example.

Figure 5: Network Topology for Local Port Mirroring Example



Mirroring All Employee Traffic for Local Analysis

To configure port mirroring for all employee traffic for local analysis, perform these tasks:

CLI Quick Configuration

To quickly configure local port mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 family ethernet-switching
set interfaces ge-0/0/1 unit 0 family inet 192.1.1/24
set interfaces ge-0/0/10 unit 0 family ethernet-switching
set ethernet-switching options analyzer employee-monitor input ingress interface ge-0/0/0.0
set ethernet-switching options analyzer employee-monitor input ingress interface ge-0/0/1.0
set ethernet-switching options analyzer employee-monitor output interface ge-0/0/10.0
```

Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the analyzer output interface:

1. Configure each interface connected to employee computers as an input interface for the port-mirror analyzer that we are calling **employee-monitor**:


```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
```
2. Configure the output analyzer interface for the **employee-monitor** analyzer. This will be the destination interface for the mirrored packets:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor output (Port Mirroring) interface ge-0/0/10.0
```

Results Check the results of the configuration:

```
[edit]
user@switch# show
ethernet-switching-options {
  analyzer employee-monitor {
    input {
```

```

        ingress {
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
        }
    }
    output {
        interface {
            ge-0/0/10.0;
        }
    }
}
}

```

Mirroring Employee-to-Web Traffic for Local Analysis

To configure port mirroring for employee to web traffic, perform these tasks:

CLI Quick Configuration

To quickly configure local port mirroring of traffic from the two ports connected to employee computers, filtering so that only traffic to the external Web is mirrored, copy the following commands and paste them into the switch terminal window:

```

[edit]
set ethernet-switching-options analyzer employee-web-monitor output interface ge-0/0/10.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then analyzer
employee-web-monitor
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

Step-by-Step Procedure

To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the local analyzer interface:

```

[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching

```

2. Configure the **employee-web-monitor** analyzer output (the input to the analyzer comes from the action of the filter):

```

[edit ethernet-switching-options]
user@switch# set analyzer employee-web-monitor output interface ge-0/0/10.0

```

3. Configure a firewall filter called **watch-employee** to send mirrored copies of employee requests to the Web to the **employee-web-monitor** analyzer. Accept all traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**). Send mirrored copies of all packets destined for the Internet (**destination port 80**) to the **employee-web-monitor** analyzer.

```

[edit firewall family ethernet-switching]
user@switch# set filter (Firewall Filters) watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28

```

```

user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port 80
user@switch# set filter watch-employee term employee-to-web then analyzer
employee-web-monitor

```

4. Apply the **watch-employee** filter to the appropriate ports:

```

[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input (Port Mirroring)
watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

Results Check the results of the configuration:

```

[edit]
user@switch# show
ethernet-switching-options {
  analyzer employee-web-monitor {
    output {
      interface ge-0/0/10.0;
    }
  }
}
...
firewall family ethernet-switching {
  filter watch-employee {
    term employee-to-corp {
      from {
        destination-address 192.0.2.16/28;
        source-address 192.0.2.16/28;
      }
      then accept {
      }
    }
    term employee-to-web {
      from {
        destination-port 80;
      }
      then analyzer employee-web-monitor;
    }
  }
}
...
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan members [employee-vlan, voice-vlan];
        filter {
          input watch-employee;
        }
      }
    }
  }
  ge-0/0/1 {
    family ethernet-switching {
      filter {
        input watch-employee;
      }
    }
  }
}

```

```
}  
}  
}  
}
```

Verification

To confirm that the configuration is correct, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 42](#)

Verifying That the Analyzer Has Been Correctly Created

Purpose Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces, and appropriate output interface.

Action You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
user@switch> show analyzer  
Analyzer name           : employee-monitor  
Output interface        : ge-0/0/10.0  
Mirror ratio            : 1  
Loss priority           : Low  
Ingress monitored interfaces : ge-0/0/0.0  
Ingress monitored interfaces : ge-0/0/1.0  
Egress monitored interfaces : None
```

Meaning This output shows that the **employee-monitor** analyzer has a ratio of 1 (mirroring every packet, the default setting), a loss priority of low (set this option to high only when the analyzer output is to a VLAN), is mirroring the traffic entering the **ge-0/0/0** and **ge-0/0/1** interfaces, and sending the mirrored traffic to the **ge-0/0/10** interface.

Related Documentation

- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
- [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
- [Understanding Port Mirroring on EX Series Switches on page 9](#)

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port

- Packets entering a VLAN on Juniper Networks EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, or EX6200 Ethernet Switches
- Packets exiting a VLAN on Juniper Networks EX8200 Ethernet Switches

You can analyze the mirrored traffic using a protocol analyzer application running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN.

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to the **remote-monitor** VLAN so that you can perform analysis from a remote monitoring station. The first example shows how to mirror all traffic entering the ports connected to employee computers. The second example shows the same scenario but includes a filter to mirror only the employee traffic going to the Web.



BEST PRACTICE: Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured port mirroring analyzers when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by:
 - Using statistical sampling.
 - Setting ratios to select statistical samples.
 - Using firewall filters.

This example describes how to configure remote port mirroring:

- [Requirements on page 43](#)
- [Overview and Topology on page 44](#)
- [Mirroring All Employee Traffic for Remote Analysis on page 44](#)
- [Mirroring Employee-to-Web Traffic for Remote Analysis on page 47](#)
- [Verification on page 51](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later for EX Series switches
- EX Series switch connected to another EX Series switch

Before you configure remote port mirroring, be sure that:

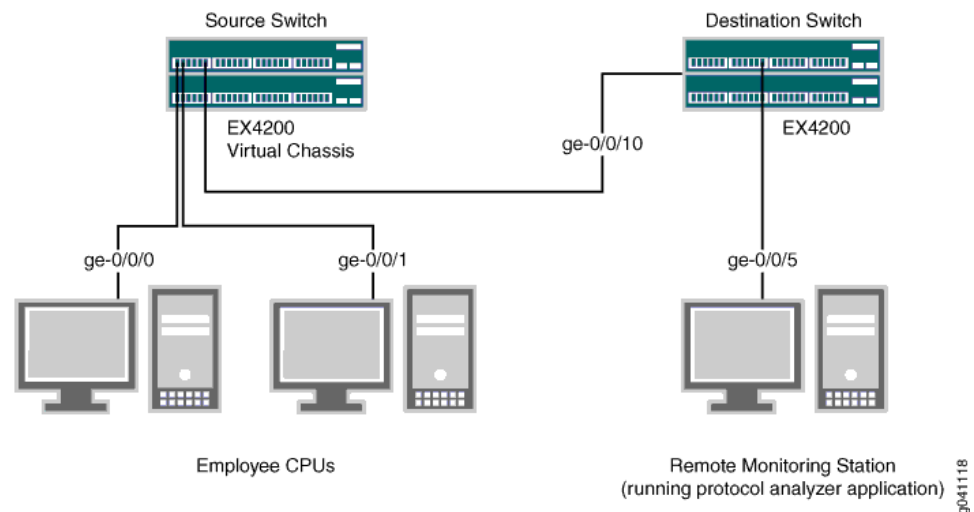
- You have an understanding of port-mirroring concepts.
- The interfaces that the analyzer will use as input interfaces have been configured on the switch.

Overview and Topology

This topic includes two related examples that describe how to configure port mirroring to the **remote-monitor** VLAN so that analysis can be performed from a remote monitoring station. The first example shows how to configure a switch to mirror all traffic from employee computers. The second example shows the same scenario, but the setup includes a filter to mirror only the employee traffic going to the Web.

Figure 6 on page 44 shows the network topology for both these example scenarios.

Figure 6: Remote Port Mirroring Example Network Topology



In this example:

- Interface **ge-0/0/0** is a Layer 2 interface, and interface **ge-0/0/1** is a Layer 3 interface (both interfaces on the source switch) that serve as connections for employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects the source switch to the destination switch.
- Interface **ge-0/0/5** is a Layer 2 interface that connects the destination switch to the remote monitoring station.
- VLAN **remote-monitor** is configured on all switches in the topology to carry the mirrored traffic.

Mirroring All Employee Traffic for Remote Analysis

To configure port mirroring for remote traffic analysis for all incoming and outgoing employee traffic, perform these tasks:

CLI Quick Configuration

To quickly configure port mirroring for remote traffic analysis for incoming and outgoing employee traffic, copy the following commands and paste them into the switch terminal window:

- Copy and paste the following commands in the source switch terminal window:

```
[edit]
set vlans remote-monitor vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set vlans remote-monitor interface ge-0/0/10 egress
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor loss-priority high
set ethernet-switching-options analyzer employee-monitor output vlan remote-monitor
```

- Copy and paste the following commands in the destination switch terminal window:

```
[edit]
set vlans remote-monitor vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set interfaces ge-0/0/5 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members 999
```

Step-by-Step Procedure

To configure basic remote port mirroring:

1. On the source switch:

- Configure the VLAN tag ID for the **remote-monitor** VLAN:

```
[edit vlans]
user@switch# set remote-monitor vlan-id 999
```

- Configure the interface on the network port connected to the destination switch for trunk mode and associate it with the **remote-monitor** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the **ge-0/0/10** interface for egress-only traffic so that traffic can only egress from the interface:

```
[edit vlans]
user@switch# set remote-monitor interface ge-0/0/10 egress
```

- Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor loss-priority high
user@switch# set analyzer employee-monitor output vlan remote-monitor
```

2. On the destination switch:

- Configure the VLAN tag ID for the **remote-monitor** VLAN:

```
[edit vlans]
user@switch# set remote-monitor vlan-id 999
```

- Configure the interface on the destination switch for trunk mode and associate it with the **remote-monitor** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the interface connected to the destination switch for trunk mode and associate it with the **remote-monitor** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/5 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/5 unit 0 family ethernet-switching vlan members 999
```

Results Check the results of the configuration on the source switch:

```
[edit]
user@switch# show
ethernet-switching-options {
  analyzer employee-monitor {
    loss-priority high;
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
      egress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
  }
  output {
    vlan {
      remote-monitor;
    }
  }
}
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
}
vlangs {
  remote-monitor {
    vlan-id 999;
    interface {
      ge-0/0/10.0
      egress;
    }
  }
}
```

Check the results of the configuration on the destination switch:

```
[edit]
user@switch# show
interfaces {
  ge0/0/5 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
}
vlangs {
  remote-monitor {
    vlan-id 999;
  }
}
```

Mirroring Employee-to-Web Traffic for Remote Analysis

To configure port mirroring for remote traffic analysis of employee to web traffic, perform these tasks:

CLI Quick Configuration To quickly configure port mirroring to mirror employee traffic to the external Web, copy the following commands and paste them into the switch terminal window:

- Copy and paste the following commands in the source switch terminal window:

```
[edit]
set ethernet-switching-options analyzer employee-web-monitor loss-priority high output vlan
999
set vlans remote-monitor vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then
accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
```

```

set firewall family ethernet-switching filter watch-employee term employee-to-web then
analyzer employee-web-monitor
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

- Copy and paste the following commands in the destination switch terminal window:

```

[edit]
set vlans remote-monitor vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set interfaces ge-0/0/5 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members 999

```

Step-by-Step Procedure To configure port mirroring of all traffic from the two ports connected to employee computers to the **remote-monitor** VLAN for use from a remote monitoring station:

1. On the source switch:

- Configure the **employee-web-monitor** analyzer:

```

[edit ethernet-switching-options]
user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
user@switch# set analyzer employee-web-monitor loss-priority high output vlan 999

```

- Configure the VLAN tag ID for the **remote-monitor** VLAN:

```

[edit vlans]
user@switch# set remote-monitor vlan-id 999

```

- Configure the interface to associate it with the **remote-monitor** VLAN:

```

[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999

```

- Configure the firewall filter called **watch-employee**:

```

[edit firewall family ethernet-switching]
user@switch# set filter (Firewall Filters) watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port
80
user@switch# set filter watch-employee term employee-to-web then analyzer
employee-web-monitor

```

- Apply the firewall filter to the employee interfaces:

```

[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input (Port Mirroring)
watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee

```

2. On the destination switch:

- Configure the VLAN tag ID for the **remote-monitor** VLAN:

```

[edit vlans]
user@switch# set remote-monitor vlan-id 999

```

- Configure the interface on the destination switch for trunk mode and associate it with the **remote-monitor** VLAN:

```

[edit interfaces]

```

```
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the interface connected to the destination switch for trunk mode and associate it with the **remote-monitor** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/5 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/5 unit 0 family ethernet-switching vlan members 999
```

Results Check the results of the configuration on the source switch:

```
[edit]
user@switch# show
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members remote-monitor;
        }
      }
    }
  }
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
}
firewall {
  family ethernet-switching {
    filter watch-employee {
      term employee-to-corp {
        from {
          source-address {
            192.0.2.16/28;
          }
          destination-address {
            192.0.2.16/28;
          }
        }
      }
    }
  }
  then accept;
```

```
    }
    term employee-to-web {
      from {
        destination-port 80;
      }
      then analyzer employee-web-monitor;
    }
  }
}
ethernet-switching-options {
  analyzer employee-web-monitor {
    loss-priority high;
    output {
      vlan {
        999;
      }
    }
  }
}
vpls {
  remote-monitor {
    vlan-id 999;
  }
}
```

Check the results of the configuration on the destination switch:

```
[edit]
user@switch# show
vpls {
  remote-monitor {
    vlan-id 999;
  }
}
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
  ge-0/0/5 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
}
```


Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 51](#)

Verifying That the Analyzer Has Been Correctly Created

Purpose	Verify that the analyzer named employee-monitor or employee-web-monitor has been created on the switch with the appropriate input interfaces and appropriate output interface.
Action	<p>You can verify the analyzer is configured as expected by using the show analyzer command. To view previously created analyzers that are disabled, go to the J-Web interface.</p> <p>To verify that the analyzer is configured as expected while monitoring all employee traffic on the source switch, run the show analyzer command on the source switch. The following output is displayed for this example configuration:</p> <pre> user@switch> show analyzer Analyzer name : employee-monitor Output VLAN : remote-monitor Mirror ratio : 1 Loss priority : High Ingress monitored interfaces : ge-0/0/0.0 Egress monitored interfaces : ge-0/0/1.0 </pre>
Meaning	This output shows that the employee-monitor analyzer has a ratio of 1 (mirroring every packet, the default), has a loss priority of high (this indicates that mirrored packets are dropped on priority in case of a congestion; the default loss-priority value is low), is mirroring the traffic entering or exiting ge-0/0/0 and ge-0/0/1 to the remote vlan called remote-monitor .
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37 • Example: Configuring Port Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX Series Switches on page 51 • Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 78 • Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 81 • Understanding Port Mirroring on EX Series Switches on page 9

Example: Configuring Port Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX Series Switches

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, or EX6200 switches
- Packets exiting a VLAN on EX8200 switches



NOTE: The feature of mirroring traffic to multiple VLAN interfaces is available only on EX8200 switches.

You can analyze the mirrored traffic using a protocol analyzer application running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN.

This example describes how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN so that you can perform analysis from a remote monitoring station. The **remote-analyzer** VLAN in this example contains multiple member interfaces. Therefore, the same traffic is mirrored to all member interfaces of the **remote-analyzer** VLAN so that mirrored packets can be sent to different remote monitoring stations. You can install applications, such as sniffers and intrusion detection systems, on remote monitoring stations to analyze these mirrored packets and to obtain useful statistical data. For instance, if there are two remote monitoring stations, you can install a sniffer on one remote monitoring station and an intrusion detection system on the other station. You can use a firewall filter analyzer configuration to forward a specific type of traffic to a remote monitoring station.



BEST PRACTICE: Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured port mirroring analyzers when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by:
 - Using statistical sampling.
 - Setting ratios to select statistical samples.
 - Using firewall filters.

This example describes how to configure remote port mirroring to multiple interfaces on an analyzer VLAN:

- [Requirements on page 53](#)
- [Overview and Topology on page 53](#)

- [Mirroring All Employee Traffic to Multiple VLAN Member Interfaces for Remote Analysis on page 55](#)
- [Verification on page 60](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later for EX Series switches
- Three EX8200 switches

Before you configure remote port mirroring, be sure that:


- You have an understanding of port-mirroring concepts.
- The interfaces that the analyzer will use as input interfaces have been configured on the switch.

Overview and Topology

This example describes how to configure port mirroring to multiple interfaces in the **remote-analyzer** VLAN so that traffic is sent to different remote monitoring stations for analysis.

[Figure 7 on page 54](#) shows the network topology for this example.

Figure 7: Remote Port Mirroring Example Network Topology Using Multiple VLAN Member Interfaces



g041119

In this example:

- Interface **ge-0/0/0** is a Layer 2 interface, and interface **ge-0/0/1** is a Layer 3 interface (both interfaces on the source switch) that serve as connections for employee computers.
- Interfaces **ge-0/0/10** and **ge-0/0/11** are Layer 2 interfaces that are connected to different destination switches.

- Interface **ge-0/0/12** is a Layer 2 interface that connects the Destination 1 switch to the remote monitoring station.
- Interface **ge-0/0/13** is a Layer 2 interface that connects the Destination 2 switch to the remote monitoring station.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.

Mirroring All Employee Traffic to Multiple VLAN Member Interfaces for Remote Analysis

To configure port mirroring to multiple VLAN member interfaces for remote traffic analysis for all incoming and outgoing employee traffic, perform these tasks:

CLI Quick Configuration

To quickly configure port mirroring for remote traffic analysis for incoming and outgoing employee traffic, copy the following commands and paste them into the switch terminal window:

- In the source switch terminal window, copy and paste the following commands:


```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set vlans remote-analyzer interface ge-0/0/10 egress
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members 999
set vlans remote-analyzer interface ge-0/0/11 egress
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor loss-priority high
set ethernet-switching-options analyzer employee-monitor output vlan remote-analyzer
```
- In the Destination 1 switch terminal window, copy and paste the following commands:


```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set vlans remote-analyzer interface ge-0/0/10 ingress
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode trunk
set ethernet-switching-options analyzer employee-monitor input ingress vlan remote-analyzer
set ethernet-switching-options analyzer employee-monitor loss-priority high output interface ge-0/0/12.0
```
- In the Destination 2 switch terminal window, copy and paste the following commands:


```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set vlans remote-analyzer interface ge-0/0/11 ingress
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
set ethernet-switching-options analyzer employee-monitor input ingress vlan remote-analyzer
set ethernet-switching-options analyzer employee-monitor loss-priority high output interface ge-0/0/13.0
```

Step-by-Step Procedure

To configure basic remote port mirroring to two VLAN member interfaces:

1. On the source switch:

- Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interfaces on the network port connected to destination switches for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/11 unit 0 family ethernet-switching vlan members 999
```

- Configure **ge-0/0/10** and **ge-0/0/11** for egress-only traffic so that traffic can only exit from the interface:

```
[edit vlans]
user@switch# set remote-analyzer interface ge-0/0/10 egress
user@switch# set remote-analyzer interface ge-0/0/11 egress
```

- Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor loss-priority high
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```

2. On the Destination 1 switch:

- Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the **ge-0/0/10** interface on the Destination 1 switch for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface for ingress traffic only:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set vlans remote-analyzer interface ge-0/0/10 ingress
```

- Configure the interface connected to the remote monitoring station for trunk mode:

```
[edit interfaces]
user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode trunk
```

- Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress vlan remote-analyzer
user@switch# set analyzer employee-monitor loss-priority high output interface ge-0/0/12.0
```

3. On the Destination 2 switch:

- Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the **ge-0/0/11** interface on the Destination 2 switch for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface only for ingress traffic:

```
[edit interfaces]
user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
user@switch# set vlans remote-analyzer interface ge-0/0/11 ingress
```

- Configure the interface connected to the remote monitoring station for trunk mode:

```
[edit interfaces]
user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
```

- Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress vlan remote-analyzer
user@switch# set analyzer employee-monitor loss-priority high output interface
ge-0/0/13.0
```

Results Check the results of the configuration on the source switch:

```
[edit]
user@switch# show
ethernet-switching-options {
  analyzer employee-monitor {
    loss-priority high;
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
      egress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
    output {
      vlan {
        remote-analyzer;
      }
    }
  }
}
vllans {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/10.0 {
        egress;
      }
      ge-0/0/11.0 {
        egress;
      }
    }
  }
}
```

```
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 999;
        }
      }
    }
  }
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          member 999;
        }
      }
    }
  }
}
```

Check the results of the configuration on the Destination 1 switch:

```
[edit]
user@switch# show
vpls {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/10.0 {
        ingress;
      }
    }
  }
}
interfaces {
  ge-0/0/10 {
    unit 0 {
      ethernet-switching {
        port-mode trunk;
      }
    }
  }
  ge-0/0/12 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}
ethernet-switching-options {
  analyzer employee-monitor {
    loss-priority high;
  }
}
```



```

    input {
      ingress {
        vlan remote-analyzer;
      }
    }
    output {
      interface {
        ge-0/0/12.0;
      }
    }
  }
}

```

Check the results of the configuration on the Destination 2 switch:

```

[edit]
user@switch# show
vpls {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/11.0 {
        ingress;
      }
    }
  }
}
interfaces {
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}
ethernet-switching-options {
  employee-monitor {
    loss-priority high;
    input {
      ingress {
        vlan remote-analyzer;
      }
    }
    output {
      interface {
        ge-0/0/13.0;
      }
    }
  }
}

```

```
}  
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 60](#)

Verifying That the Analyzer Has Been Correctly Created

Purpose	Verify that the analyzer named employee-monitor has been created on the switch with the appropriate input interfaces and appropriate output interface.
Action	<p>You can verify the analyzer is configured as expected by using the show analyzer command. To view previously created analyzers that are disabled, go to the J-Web interface.</p> <p>To verify that the analyzer is configured as expected while monitoring all employee traffic on the source switch, run the show analyzer command on the source switch. The following output is displayed for this example configuration:</p> <pre>user@switch> show analyzer Analyzer name : employee-monitor Output VLAN : remote-analyzer Mirror ratio : 1 Loss priority : High Ingress monitored interfaces : ge-0/0/0.0 Ingress monitored interfaces : ge-0/0/1.0</pre>
Meaning	This output shows that the employee-monitor analyzer has a ratio of 1 (mirroring every packet, the default), has a loss priority of high (set this option to high whenever the analyzer output is to a VLAN), is mirroring the traffic entering ge-0/0/0 and ge-0/0/1 , and is sending the mirrored traffic to the analyzer called remote-analyzer .
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42• Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 78• Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 81• Understanding Port Mirroring on EX Series Switches on page 9

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, or EX6200 switches
- Packets exiting a VLAN on EX8200 switches

You can analyze the mirrored traffic using a protocol analyzer application running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN.

This topic includes an example that describes how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN through a transit switch, so that you can perform analysis from a remote monitoring station.



BEST PRACTICE: Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured port mirroring analyzers when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by:
 - Using statistical sampling.
 - Setting ratios to select statistical samples.
 - Using firewall filters.

This example describes how to configure remote port mirroring through a transit switch:

- [Requirements on page 61](#)
- [Overview and Topology on page 62](#)
- [Mirroring All Employee Traffic for Remote Analysis Through a Transit Switch on page 63](#)
- [Verification on page 67](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.5 or later for EX Series switches

- EX3200 or EX4200 switch connected to another EX3200 or EX4200 switch through a third EX3200 or EX4200 switch

Before you configure remote port mirroring, be sure that:

- You have an understanding of port-mirroring concepts.
- The interfaces that the analyzer will use as input interfaces have been configured on the switch.

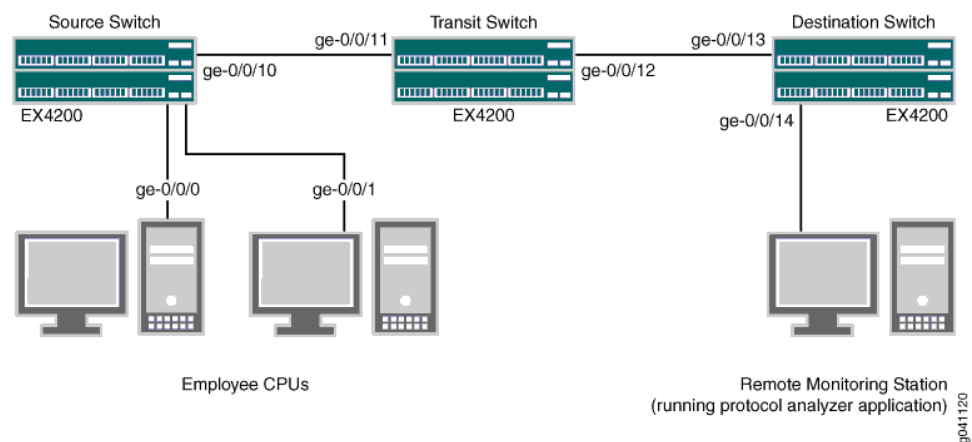
Overview and Topology

This example describes how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN through a transit switch so that you can perform analysis from a remote monitoring station. The example shows how to configure a switch to mirror all traffic from employee computers to a remote analyzer.

In this configuration, an analyzer session is required on the destination switch to mirror incoming traffic from the analyzer VLAN to the egress interface to which the remote monitoring station is connected. You must disable MAC learning on the transit switch for the **remote-analyzer** VLAN so that MAC learning is disabled for all member interfaces of the **remote-analyzer** VLAN on the transit switch.

Figure 8 on page 62 shows the network topology for this example.

Figure 8: Remote Port Mirroring Example Through a Transit Switch
Network Topology



In this example:

- Interface **ge-0/0/0** is a Layer 2 interface, and interface **ge-0/0/1** is a Layer 3 interface (both interfaces on the source switch) that serve as connections for employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects to the transit switch.
- Interface **ge-0/0/11** is a Layer 2 interface on the transit switch.
- Interface **ge-0/0/12** is a Layer 2 interface on the transit switch and connects to the destination switch.

- Interface **ge-0/0/13** is a Layer 2 interface on the destination switch .
- Interface **ge-0/0/14** is a Layer 2 interface on the destination switch and connects to the remote monitoring station.
- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.

Mirroring All Employee Traffic for Remote Analysis Through a Transit Switch

To configure port mirroring for remote traffic analysis through a transit switch, for all incoming and outgoing employee traffic, perform these tasks:

CLI Quick Configuration

To quickly configure port mirroring for remote traffic analysis through a transit switch, for incoming and outgoing employee traffic, copy the following commands and paste them into the switch terminal window:

- Copy and paste the following commands in the source switch (monitored switch) terminal window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor output vlan remote-analyzer
```

- Copy and paste the following commands in the transit switch window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set vlans remote-analyzer interface ge-0/0/11 ingress
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode trunk
set vlans remote-analyzer interface ge-0/0/12 egress
set vlans remote-analyzer no-mac-learning
```

- Copy and paste the following commands in the destination switch window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
set vlans remote-analyzer interface ge-0/0/13 ingress
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set ethernet-switching-options analyzer employee-monitor input ingress vlan remote-analyzer
set ethernet-switching-options analyzer employee-monitor loss-priority high output interface ge-0/0/14.0
```

Step-by-Step Procedure

To configure remote port mirroring through a transit switch:

1. On the source switch:
 - Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the interfaces on the network port connected to transit switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

- Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/0.0
user@switch# set analyzer employee-monitor input egress interface ge-0/0/1.0
user@switch# set analyzer employee-monitor loss-priority high
user@switch# set analyzer employee-monitor output vlan remote-analyzer
```

2. On the transit switch:

- Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the **ge-0/0/11** interface for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface for ingress traffic only:

```
[edit interfaces]
user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
user@switch# set vlans remote-analyzer interface ge-0/0/11 ingress
```

- Configure the **ge-0/0/12** interface for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface for egress traffic only:

```
[edit interfaces]
user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode trunk
user@switch# set vlans remote-analyzer interface ge-0/0/12 egress
```

- Configure the **no-mac-learning** option for the **remote-analyzer** VLAN to disable MAC learning on all interfaces that are members of the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set vlans remote-analyzer no-mac-learning
```

3. On the destination switch:

- Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

- Configure the **ge-0/0/13** interface for trunk mode, associate it with the **remote-analyzer** VLAN, and set the interface for ingress traffic only:

```
[edit interfaces]
user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch# set vlans remote-analyzer interface ge-0/0/13 ingress
```

- Configure the interface connected to the remote monitoring station for trunk mode:

```
[edit interfaces]
user@switch# set ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
```

- Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
```

```

user@switch# set analyzer employee-monitor input ingress vlan remote-analyzer
user@switch# set analyzer employee-monitor loss-priority high output interface
ge-0/0/14.0

```

Results Check the results of the configuration on the source switch:

```

[edit]
user@switch# show
ethernet-switching-options {
  analyzer employee-monitor {
    input {
      ingress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
      egress {
        interface ge-0/0/0.0;
        interface ge-0/0/1.0;
      }
    }
    output {
      vlan {
        remote-analyzer;
      }
    }
  }
}
vlangs {
  remote-analyzer {
    vlan-id 999;
  }
}
interfaces {
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          member 999;
        }
      }
    }
  }
}
}

```

Check the results of the configuration on the transit switch:

```

[edit]
user@switch# show
vlangs {
  remote-analyzer {
    vlan-id 999;
    interface {
      ge-0/0/11.0 {
        ingress;
      }
    }
  }
}

```

```
        ge-0/0/12.0 {
            egress;
        }
    }
    no-mac-learning;
}
}
interfaces {
    ge-0/0/11 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
    ge-0/0/12 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
}
}
```

Check the results of the configuration on the destination switch:

```
[edit]
user@switch# show
vpls {
    remote-analyzer {
        vlan-id 999;
        interface {
            ge-0/0/13.0 {
                ingress;
            }
        }
    }
}
interfaces {
    ge-0/0/13 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
    ge-0/0/14 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
}
}
ethernet-switching-options {
    analyzer employee-monitor {
```



```

    loss-priority high;
    input {
        ingress {
            vlan remote-analyzer;
        }
    }
    output {
        interface {
            ge-0/0/14.0;
        }
    }
}

```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That the Analyzer Has Been Correctly Created on page 67](#)

Verifying That the Analyzer Has Been Correctly Created

Purpose	Verify that the analyzer named employee-monitor has been created on the switch with the appropriate input interfaces and the appropriate output interface.
Action	<p>You can verify the analyzer is configured as expected by using the show analyzer command. To view previously created analyzers that are disabled, go to the J-Web interface.</p> <p>To verify that the analyzer is configured as expected while monitoring all employee traffic on the source switch, run the show analyzer command on the source switch. The following output is displayed for this example configuration:</p> <pre> user@switch> show analyzer Analyzer name : employee-monitor Output VLAN : remote-analyzer Mirror ratio : 1 Loss priority : High Ingress monitored interfaces : ge-0/0/0.0 Ingress monitored interfaces : ge-0/0/1.0 </pre>
Meaning	This output shows that the employee-monitor analyzer has a ratio of 1 (mirroring every packet, the default), has a loss priority of high (set this option to high whenever the analyzer output is to a VLAN), is mirroring the traffic entering ge-0/0/0 and ge-0/0/1 , and is sending the mirrored traffic to the analyzer called remote-analyzer .
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42 • Example: Configuring Port Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX Series Switches on page 51 • Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37

- [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
- [Understanding Port Mirroring on EX Series Switches on page 9](#)

Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches

Ethernet interfaces on EX Series switches and Junos OS for EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM).

This example describes how to enable and configure OAM CFM on a Gigabit Ethernet interface:

- [Requirements on page 68](#)
- [Overview and Topology on page 68](#)
- [Configuring Ethernet OAM Connectivity Fault Management on Switch 1 on page 68](#)
- [Configuring Ethernet OAM Connectivity Fault Management on Switch 2 on page 69](#)
- [Verification on page 71](#)

Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.2 or later for EX Series switches
- Two EX Series switches connected by a point-to-point Gigabit Ethernet link

Overview and Topology

CFM can be used to monitor the physical link between two switches. In the following example, two switches are connected by a point-to-point Gigabit Ethernet link. The link between these two switches is monitored using CFM.

Configuring Ethernet OAM Connectivity Fault Management on Switch 1

CLI Quick Configuration	<p>To quickly configure Ethernet OAM CFM, copy the following commands and paste them into the switch terminal window:</p> <pre>[edit protocols oam ethernet connectivity-fault-management maintenance-domain] set name-format character-string set maintenance-domain private level 0 set maintenance-association private-ma set continuity-check hold-interval 1s</pre>
--------------------------------	--

Step-by-Step Procedure

To enable and configure OAM CFM on switch 1:

1. Specify the maintenance domain name format:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]
user@switch1# set name-format character-string
```
2. Specify the maintenance domain name and the maintenance domain level:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch1# set maintenance-domain private level 0
```
3. Create a maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain private]
user@switch1# set maintenance-association private-ma
```
4. Enable the continuity check protocol and specify the continuity check hold interval:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain private maintenance-association private-ma]
user@switch1# set continuity-check hold-interval 1s
```
5. Configure the maintenance association end point (MEP):

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain private maintenance-association private-ma]
user@switch1# set mep 100 interface ge-1/0/1 auto-discovery direction down
```

Results

Check the results of the configuration.

```
[edit]
user@switch1 > show

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain private {
          level 0;
          maintenance-association private-ma {
            continuity-check {
              interval 1s;
            }
            mep 100 {
              interface ge-1/0/1;
              auto-discovery;
              direction down;
            }
          }
        }
      }
    }
  }
}
```

Configuring Ethernet OAM Connectivity Fault Management on Switch 2**CLI Quick Configuration**

To quickly configure Ethernet OAM CFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]
```

```
set name-format character-string
set maintenance-domain private level 0
set maintenance-association private-ma
set continuity-check hold-interval 1s
```

Step-by-Step Procedure

The configuration on switch 2 mirrors that on switch 2.

1. Specify the maintenance domain name format:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch2# set name-format character-string
```
2. Specify the maintenance domain name and the maintenance domain level:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch2# set maintenance-domain private level 0
```
3. Create a maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
private]
user@switch2# set maintenance-association private-ma
```
4. Enable the continuity check protocol and specify the continuity check hold interval:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
private maintenance-association private-ma]
user@switch2# set continuity-check hold-interval 1s
```
5. Configure the maintenance association end point (MEP)

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
private maintenance-association private-ma]
user@switch2# set mep 200 interface ge-0/2/5 auto-discovery direction down
```

Results

Check the results of the configuration.

```
[edit]
user@switch2 > show

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain private {
          level 0;
          maintenance-association private-ma {
            continuity-check {
              interval 1s;
            }
            mep 200 {
              interface ge-0/2/5;
              auto-discovery;
              direction down;
            }
          }
        }
      }
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying That OAM CFM Has Been Configured Properly on page 71](#)

Verifying That OAM CFM Has Been Configured Properly

- Purpose** Verify that OAM CFM has been configured properly.
- Action** Use the `show oam ethernet connectivity-fault-management interfaces detail` command:
- ```
user@switch1# show oam ethernet connectivity-fault-management interfaces detail
```

## Sample Output

```
Interface name: ge-1/0/1.0, Interface status: Active, Link status: Up
Maintenance domain name: private, Format: string, Level: 0
Maintenance association name: private-ma, Format: string
Continuity-check status: enabled, Interval: 1ms, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
 Remote MEP not receiving CCM : no
 Erroneous CCM received : yes
 Cross-connect CCM received : no
 RDI sent by some MEP : yes
Statistics:
 CCMs sent : 76
 CCMs received out of sequence : 0
 LBMs sent : 0
 Valid in-order LBRs received : 0
 Valid out-of-order LBRs received : 0
 LBRs received with corrupted data : 0
 LBRs sent : 0
 LTMs sent : 0
 LTMs received : 0
 LTRs sent : 0
 LTRs received : 0
 Sequence number of next LTM request : 0
Remote MEP count: 2
 Identifier MAC address State Interface
 2001 00:90:69:0b:7f:71 ok ge-0/2/5.0
```

- Meaning** When the output displays that continuity-check status is **enabled** and displays details of the remote MEP, it means that connectivity fault management (CFM) has been configured properly.

- Related Documentation**
- [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch on page 22](#)
  - [Junos OS Network Interfaces Configuration Guide](#)

## Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet interface:

- [Requirements on page 72](#)
- [Overview and Topology on page 72](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 1 on page 72](#)
- [Configuring Ethernet OAM Link Fault Management on Switch 2 on page 73](#)
- [Verification on page 74](#)

### Requirements

This example uses the following hardware and software components:

- Junos OS Release 9.4 or later for EX Series switches
- Two EX3200 or EX4200 switches connected directly

### Overview and Topology

Junos OS for EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two EX4200 switches connected directly. Before you begin configuring Ethernet OAM LFM on two switches, connect the two switches directly through a trunk interface.

### Configuring Ethernet OAM Link Fault Management on Switch 1

**CLI Quick Configuration** To quickly configure Ethernet OAM LFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/0
set interface ge-0/0/0 link-discovery active
set interface ge-0/0/0 pdu-interval 800
set interface ge-0/0/0 remote-loopback
```

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step-by-Step Procedure</b> | <p>To configure Ethernet OAM LFM on switch 1:</p> <ol style="list-style-type: none"> <li>1. Enable IEEE 802.3ah OAM support on an interface:           <pre>[edit protocols oam ethernet link-fault-management] user@switch1# set interface (OAM LFM) ge-0/0/0</pre> </li> <li>2. Specify that the interface initiates the discovery process by configuring the link discovery mode to <b>active</b>:           <pre>[edit protocols oam ethernet link-fault-management] user@switch1# set interface ge-0/0/0 link-discovery active</pre> </li> <li>3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800 on switch 1:           <pre>[edit protocols oam ethernet link-fault-management] user@switch1# set interface pdu-interval 800</pre> </li> <li>4. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Ensure that the remote DTE supports remote loopback mode. To set the remote DTE in loopback mode           <pre>[edit protocols oam ethernet link-fault-management] user@switch1# set interface ge-0/0/0.0 remote-loopback</pre> </li> </ol> |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### Results

Check the results of the configuration:

```
[edit]
user@switch1# show

protocols {
 oam {
 ethernet {
 link-fault-management {
 interface ge-0/0/0 {
 pdu-interval 800;
 link-discovery active;
 remote-loopback;
 }
 }
 }
 }
}
```

## Configuring Ethernet OAM Link Fault Management on Switch 2

|                                |                                                                                                                                |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <b>CLI Quick Configuration</b> | To quickly configure Ethernet OAM LFM on switch 2, copy the following commands and paste them into the switch terminal window: |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------|

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/1
set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

|                               |                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Step-by-Step Procedure</b> | <p>To configure Ethernet OAM LFM on switch 2:</p> <ol style="list-style-type: none"> <li>1. Enable OAM on the peer interface on switch 2:           <pre>[edit protocols oam ethernet link-fault-management] user@switch2# set interface ge-0/0/1</pre> </li> <li>2. Enable remote loopback support for the local interface:           <pre>[edit protocols oam ethernet link-fault-management]</pre> </li> </ol> |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

```
user@switch2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

**Results** Check the results of the configuration:

```
[edit]
```

```
user@switch2# show
```

```
protocols {
 oam {
 ethernet {
 link-fault-management {
 interface ge-0/0/1 {
 negotiation-options {
 allow-remote-loopback;
 }
 }
 }
 }
 }
}
```

## Verification

---

### Verifying That OAM LFM Has Been Configured Properly

**Purpose** Verify that OAM LFM has been configured properly.

**Action** Use the `show oam ethernet link-fault-management` command:

```
user@switch1#show oam ethernet link-fault-management
```

## Sample Output

```
Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported
```

**Meaning** When the output displays the MAC address and the discover state is **Send Any**, it means that OAM LFM has been configured properly.

**Related Documentation**

- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 97](#)
- [Understanding Ethernet OAM Link Fault Management on page 21](#)



## CHAPTER 7

# Configuration Tasks

- [Configuring SNMP \(J-Web Procedure\) on page 75](#)
- [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 84](#)
- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 85](#)
- [Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch \(CLI Procedure\) on page 92](#)
- [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 93](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 97](#)
- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 99](#)
- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)
- [Configuring an Iterator Profile on a Switch \(CLI Procedure\) on page 102](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 103](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)

## Configuring SNMP (J-Web Procedure)

---



**NOTE:** This topic applies only to the J-Web Application package.

You can use the J-Web interface to define system identification information, create SNMP communities, create SNMP trap groups, and configure health monitor options for EX Series switches.

To configure SNMP features:

1. Select **Configure > Services > SNMP**.

2. Enter information into the configuration page for SNMP as described in [Table 5 on page 76](#).
3. To apply the configuration click **Apply**.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

**Table 5: SNMP Configuration Page**

| Field                                   | Function                                                                                                                                                                                                                                                                                                                                                                | Your Action                                                                                   |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| <b>Identification</b>                   |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| Contact Information                     | Free-form text string that specifies an administrative contact for the system.                                                                                                                                                                                                                                                                                          | Type contact information for the administrator of the system (such as name and phone number). |
| System Description                      | Free-form text string that specifies a description for the system.                                                                                                                                                                                                                                                                                                      | Type information that describes the system                                                    |
| Local Engine ID                         | Provides an administratively unique identifier of an SNMPv3 engine for system identification.<br><br>The local engine ID contains a prefix and a suffix. The prefix is formatted according to specifications defined in RFC 3411. The suffix is defined by the local engine ID. Generally, the local engine ID suffix is the MAC address of Ethernet management port 0. | Type the MAC address of Ethernet management port 0.                                           |
| System Location                         | Free-form text string that specifies the location of the system.                                                                                                                                                                                                                                                                                                        | Type location information for the system (lab name or rack name, for example).                |
| System Override Name                    | Free-form text string that overrides the system hostname.                                                                                                                                                                                                                                                                                                               | Type the hostname of the system.                                                              |
| <b>Communities</b>                      |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| To add a community, click <b>Add</b>    |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| Community Name                          | Specifies the name of the SNMP community.                                                                                                                                                                                                                                                                                                                               | Type the name of the community being added.                                                   |
| Authorization                           | Specifies the type of authorization (either read-only or read-write) for the SNMP community being configured.                                                                                                                                                                                                                                                           | Select the authorization (either read-only or read-write) from the list.                      |
| <b>Traps</b>                            |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |
| To add a trap group, click <b>Add</b> . |                                                                                                                                                                                                                                                                                                                                                                         |                                                                                               |

Table 5: SNMP Configuration Page (*continued*)

| Field                    | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Trap Group Name          | Specifies the name of the SNMP trap group being configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Type the name of the group being added.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Categories               | Specifies which trap categories are added to the trap group being configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <ul style="list-style-type: none"> <li>To generate traps for authentication failures, select <b>Authentication</b>.</li> <li>To generate traps for chassis and environment notifications, select <b>Chassis</b>.</li> <li>To generate traps for configuration changes, select <b>Configuration</b>.</li> <li>To generate traps for link-related notifications (up-down transitions), select <b>Link</b>.</li> <li>To generate traps for remote operation notifications, select <b>Remote operations</b>.</li> <li>To generate traps for remote network monitoring (RMON), select <b>RMON alarm</b>.</li> <li>To generate traps for routing protocol notifications, select <b>Routing</b>.</li> <li>To generate traps on system warm and cold starts, select <b>Startup</b>.</li> <li>To generate traps on Virtual Router Redundancy Protocol (VRRP) events (such as new-master or authentication failures), select <b>VRRP events</b>.</li> </ul> |
| Targets                  | Specifies one or more hostnames or IP addresses for the systems to receive SNMP traps generated by the trap group being configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>Enter the hostname or IP address, in dotted decimal notation, of the target system to receive the SNMP traps.</li> <li>Click <b>Add</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Health Monitoring        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Enable Health Monitoring | <p>Enables the SNMP health monitor on the switch. The health monitor periodically (over the time you specify in the interval field) checks the following key indicators of switch health:</p> <ul style="list-style-type: none"> <li>Percentage of file storage used</li> <li>Percentage of Routing Engine CPU used</li> <li>Percentage of Routing Engine memory used</li> <li>Percentage of memory used for each system process</li> <li>Percentage of CPU used by the forwarding process</li> <li>Percentage of memory used for temporary storage by the forwarding process</li> </ul> | <p>Select the check box to enable the health monitor and configure options. Clear the check box to disable the health monitor.</p> <p><b>NOTE:</b> If you select the <b>Enable Health Monitoring</b> check box and do not specify options, then SNMP health monitoring is enabled with default values.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Interval                 | <p>Specifies the sampling frequency, in seconds, over which the key health indicators are sampled and compared with the rising and falling thresholds.</p> <p>For example, if you configure the interval as 100 seconds, the values are checked every 100 seconds.</p>                                                                                                                                                                                                                                                                                                                   | <p>Enter an interval time, in seconds, from 1 through 2147483647.</p> <p>The default value is 300 seconds (5 minutes).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 5: SNMP Configuration Page (*continued*)

| Field             | Function                                                                                                                                                                                                                                                                                                       | Your Action                                                                                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Rising Threshold  | <p>Specifies the value at which SNMP generates an event (trap and system log message) when the value of a sampled indicator is increasing.</p> <p>For example, if the rising threshold is 90 (the default), SNMP generates an event when the value of any key indicator reaches or exceeds 90 percent.</p>     | <p>Enter a value from <b>0</b> through <b>100</b>. The default value is <b>90</b>.</p>                                                                                               |
| Falling Threshold | <p>Specifies the value at which SNMP generates an event (trap and system log message) when the value of a sampled indicator is decreasing.</p> <p>For example, if the falling threshold is 80 (the default), SNMP generates an event when the value of any key indicator falls back to 80 percent or less.</p> | <p>Enter a value from <b>0</b> through <b>100</b>. The default value is <b>80</b>.</p> <p><b>NOTE:</b> The falling threshold value must be less than the rising threshold value.</p> |

- Related Documentation**
- *Monitoring System Process Information*
  - *Monitoring System Properties*

## Configuring Port Mirroring to Analyze Traffic (CLI Procedure)

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX2200, EX3200, EX3300, EX4200, EX4500, or EX6200 switches
- Packets exiting a VLAN on EX8200 switches



**BEST PRACTICE:** Mirror only necessary packets to reduce potential performance impact. We recommend that you:

- Disable your configured port mirroring analyzers when you are not using them.
- Specify individual interfaces as input to analyzers rather than specifying all interfaces as input.
- Limit the amount of mirrored traffic by:
  - Using statistical sampling.
  - Setting ratios to select statistical samples.
  - Using firewall filters.



**NOTE:** If you want to create additional analyzers without deleting the existing analyzer, first disable the existing analyzer using the `disable analyzer analyzer-name` command or the J-Web configuration page for port mirroring.



**NOTE:** Interfaces used as output for an analyzer must be configured as family `ethernet-switching`.

- [Configuring Port Mirroring for Local Traffic Analysis on page 79](#)
- [Configuring Port Mirroring for Remote Traffic Analysis on page 79](#)
- [Filtering the Traffic Entering an Analyzer on page 80](#)

## Configuring Port Mirroring for Local Traffic Analysis

To mirror interface traffic or VLAN traffic on the switch to an interface on the switch:

1. Choose a name for the analyzer—in this case, **employee-monitor**—and specify the input—in this case, packets entering **ge-0/0/0** and **ge-0/0/1**:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input (Port Mirroring) ingress interface
ge-0/0/0.0
```

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
```

2. Optionally, you can specify a statistical sampling of the packets by setting a ratio:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor ratio 200
```

When the ratio is set to 200, 1 of every 200 packets is mirrored to the analyzer. You can use statistical sampling to reduce the volume of mirrored traffic, as a high volume of mirrored traffic can be performance intensive for the switch. On EX8200 switches, you can set a ratio only for ingress packets.

3. Configure the destination interface for the mirrored packets:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor output (Port Mirroring) interface ge-0/0/10.0
```

## Configuring Port Mirroring for Remote Traffic Analysis

To mirror traffic that is traversing interfaces or a VLAN on the switch to a VLAN for analysis from a remote location:

1. Configure a VLAN to carry the mirrored traffic. This VLAN is called **remote-analyzer** and given the ID of 999 by convention in this documentation:

```
[edit]
user@switch# set vlans remote-analyzer vlan-id 999
```

2. Set the uplink module interface that is connected to the distribution switch to trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit]
```

```
user@switch# set interfaces ge-0/1/1 unit 0 family ethernet-switching port-mode trunk vlan members 999
```

3. Configure the analyzer:

- a. Choose a name and set the loss priority to high. Loss priority should always be set to high when configuring for remote port mirroring:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor loss-priority (Port Mirroring) high
```

- b. Specify the traffic to be mirrored—in this example the packets entering ports **ge-0/0/0** and **ge-0/0/1**:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/1.0
```

- c. Specify the **remote-analyzer** VLAN as the output for the analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor output (Port Mirroring) vlan 999
```

4. Optionally, you can specify a statistical sampling of the packets by setting a ratio:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor ratio 200
```

When the ratio is set to 200, 1 out of every 200 packets is mirrored to the analyzer. You can use this to reduce the volume of mirrored traffic as a very high volume of mirrored traffic can be performance intensive for the switch.

## Filtering the Traffic Entering an Analyzer

To filter which packets are mirrored to an analyzer, create the analyzer and then use it as the action in the firewall filter. You can use firewall filters in both local and remote port mirroring configurations.

If the same analyzer is used in multiple filters or terms, the packets are copied to the analyzer output port or analyzer VLAN only once.

To filter mirrored traffic, create an analyzer and then create a firewall filter. The filter can use any of the available match conditions and must have an action of **analyzer analyzer-name**. The action of the firewall filter provides the input to the analyzer.

To configure port mirroring with filters:

1. Configure the analyzer name (here, **employee-monitor**) and the output:

- a. For local analysis, set the output to the local interface to which you will connect the computer running the protocol analyzer application:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor output interface ge-0/0/10.0
```

- b. For remote analysis, set the loss priority to high and set the output to the **remote-analyzer** VLAN:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor loss-priority high output vlan 999
```

2. Create a firewall filter using any of the available match conditions and specify the action as **analyzer employee-monitor**:

This step shows a firewall filter called **example-filter**, with two terms:

- a. Create the first term to define the traffic that should not pass through to the analyzer:

```
[edit firewall family ethernet-switching]
user@switch# set filter (Firewall Filters) example-filter term no-analyzer from
source-address ip-address
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term no-analyzer from destination-address
ip-address
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term no-analyzer then accept
```

- b. Create the second term to define the traffic that should pass through to the analyzer:

```
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term to-analyzer from destination-port 80
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term to-analyzer then analyzer employee-monitor
[edit firewall family ethernet-switching]
user@switch# set filter example-filter term to-analyzer then accept
```

3. Apply the firewall filter to the interfaces or VLAN that are input to the analyzer:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input
example-filter
[edit]
user@switch# set vlan (802.1Q Tagging) remote-analyzer filter input example-filter
```

#### Related Documentation

- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
- [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37](#)
- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
- [Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on EX Series Switches](#)
- [Understanding Port Mirroring on EX Series Switches on page 9](#)
- [Firewall Filters for EX Series Switches Overview](#)

## Configuring Port Mirroring to Analyze Traffic (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on EX2200, EX3200, EX3300, EX4200, EX4300, EX4500, EX6200 switches
- Packets exiting a VLAN on EX8200 switches

To configure port mirroring on an EX Series switch using the J-Web interface:

1. Select **Configure > Security > Port Mirroring**.

The top of the screen displays analyzer details such as the name, status, analyzer port, ratio, and loss priority.

The bottom of the screen lists ingress and egress ports of the selected analyzer.



**NOTE:** After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See [Using the Commit Options to Commit Configuration Changes](#) for details about all commit options.

2. Click one of the following options:

- **Add**—Add an analyzer. Enter information as specified in [Table 6 on page 82](#).
- **Edit**—Modify details of the selected analyzer. Enter information as specified in [Table 6 on page 82](#).
- **Delete**—Delete the selected analyzer.
- **Enable/Disable**—Enable or disable the selected analyzer (toggle).



**NOTE:** On EX2200, EX3200, EX4200, and EX4500 switches, only one analyzer can be enabled at a time. On EX8200 switches, a maximum of seven analyzers can be enabled. On EX4300 switches a maximum of four Analyzers/Port Mirror instances can be enabled.



**NOTE:** When an analyzer is deleted or disabled, any filter association is removed.

**Table 6: Port Mirroring Configuration Settings**

| Field         | Function                            | Your Action                   |
|---------------|-------------------------------------|-------------------------------|
| Analyzer Name | Specifies the name of the analyzer. | Type a name for the analyzer. |



Table 6: Port Mirroring Configuration Settings (*continued*)

| Field                                                                                 | Function                                                                                                                                                                                                                                                                                                                                                        | Your Action                                                                                                                                                                                                                                                                  |
|---------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ratio<br><br><b>NOTE:</b> This option is not supported on EX4300 switches.            | Specifies the ratio of packets to be mirrored. For example: <ul style="list-style-type: none"> <li>A ratio of 1 sends copies of all packets.</li> <li>A ratio of 2047 sends copies of 1 out of every 2047 packets.</li> </ul>                                                                                                                                   | Enter a number from 0 through 2047.                                                                                                                                                                                                                                          |
| Loss Priority<br><br><b>NOTE:</b> This option is not supported on EX4300 switches.    | Specifies the loss priority of the mirrored packets.<br><br>By default, the switch applies a lower priority to mirrored data than to regular port-to-port data—mirrored traffic is dropped in preference to regular traffic when capacity is exceeded.<br><br>For port-mirroring configurations with output to an analyzer VLAN, set the loss priority to high. | Keep the default of low, unless the output is to a VLAN.                                                                                                                                                                                                                     |
| Analyzer Port                                                                         | Specifies a local interface or VLAN to which mirrored packets are sent.<br><br><b>NOTE:</b> A VLAN must have only one associated interface to be specified as an analyzer interface.                                                                                                                                                                            | Click <b>Select</b> . In the Select Analyzer Port/VLAN window, select either port or VLAN as the <b>Analyzer Type</b> . Next, select the required port or VLAN. For an EX8200 Virtual Chassis configuration, select the member, FPC, and the port (interface) from the list. |
| Analyzer Type<br><br><b>NOTE:</b> This option is supported only on EX4300 switches.   | Specifies the analyzer type.                                                                                                                                                                                                                                                                                                                                    | Select the <b>Analyzer Type</b> from the list.                                                                                                                                                                                                                               |
| No Filter check<br><br><b>NOTE:</b> This option is supported only on EX4300 switches. | Enable this option to skip checking for filters on port-mirroring instance.                                                                                                                                                                                                                                                                                     | To enable this option, select the check box.                                                                                                                                                                                                                                 |
| Ingress                                                                               | Specifies interfaces or VLANs for which entering traffic is mirrored.                                                                                                                                                                                                                                                                                           | Click <b>Add</b> . For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.<br><br>Click <b>Remove</b> to delete an ingress interface or VLAN.                                                                                  |
| Egress                                                                                | Specifies interfaces for which exiting traffic is mirrored.                                                                                                                                                                                                                                                                                                     | Click <b>Add</b> and select <b>Port</b> or <b>VLAN</b> . For an EX8200 Virtual Chassis configuration, select the member, FPC, and the interface from the list.<br><br>Click <b>Remove</b> to remove egress interfaces.                                                       |

**Related Documentation** • [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)

- [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37](#)
- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
- [Understanding Port Mirroring on EX Series Switches on page 9](#)

## Configuring sFlow Technology for Network Monitoring (CLI Procedure)

---

sFlow technology is a network monitoring technology for high-speed switched or routed networks. It is a technology that is based on statistical sampling. You can configure sFlow technology to continuously monitor traffic at wire speed on all interfaces simultaneously. Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks*.

To configure sFlow features:

1. Configure the IP address and the UDP port of the collector:

```
[edit protocols]
user@switch# set sflow collector ip-address udp-port port-number
```

The default UDP port is 6343,

2. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces interface-name
```



**NOTE:** You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a link aggregation group (LAG), but you can enable it on the member interfaces of a LAG.

3. Specify in seconds how often the sFlow agent polls interfaces:

```
[edit protocols sflow]
user@switch# set polling-interval seconds
```



**NOTE:** Specify 0 if you do not want to poll the interface.

4. Specify the rate at which packets must be sampled. You can specify either an egress or an ingress sampling rate, or both.



**NOTE:** We recommend that you configure the same sampling rates on all the ports on a line card. If you configure different sampling rates on different ports, the lowest value is used for all ports. You could still configure different rates on different line cards.

To specify an egress sampling rate:

```
[edit protocols sflow]
user@switch# set sample-rate egress number
```

To specify an ingress sampling rate:

```
[edit protocols sflow]
user@switch# set sample-rate ingress number
```

5. To configure the polling interval and the egress and ingress sampling rates at the interface level:

```
[edit protocols sflow interfaces interface-name]
user@switch# set polling-interval seconds
[edit protocols sflow interfaces]
user@switch# set sample-rate egress number
[edit protocols sflow interfaces]
user@switch# set sample-rate ingress number
```



**NOTE:** The interface-level configuration overrides the global configuration.

6. To specify an IP address to be used as the agent ID for the sFlow agent:

```
[edit protocols sflow]
user@switch# set agent-id ip-address
```

7. To specify the source IP address to be used for sFlow datagrams:

```
[edit protocols sflow]
user@switch# set source-ip ip-address
```

#### Related Documentation

- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33](#)
- [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 17](#)

## Configuring Real-Time Performance Monitoring (J-Web Procedure)



**NOTE:** This topic applies only to the J-Web Application package.

Real-time performance monitoring (RPM) in EX Series switches enables you to configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter. Jitter is the difference in relative transit time between two consecutive probes. You can set up probe owners and configure one or more performance probe tests under each probe owner.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.

- Set thresholds to trigger SNMP traps when threshold values are exceeded. You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test.
- Determine automatically whether a path exists between a host switch and its configured Border Gateway Protocol (BGP) neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets.

EX Series switches support the following tests and probe types:

- Ping tests:
  - ICMP echo
  - ICMP timestamp
- HTTP tests:
  - HTTP get (not available for BGP RPM services)
- UDP and TCP tests with user-configured ports:
  - UDP echo
  - TCP connection
  - UDP timestamp

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You must configure both the requester and the responder to timestamp the RPM packets. The RPM features provides an additional configuration option to set one-way hardware timestamps. Use one-way timestamps when you want information about one-way, rather than round-trip, times for packets to traverse the network between the requester and the responder.



NOTE:

- EX Series switches support hardware timestamps for UDP and ICMP probes. EX Series switches do not support hardware timestamps for HTTP or TCP probes.
  - If the responder does not support hardware timestamps, RPM can only report the round-trip measurements, it cannot calculate round-trip jitter.
  - In EX Series switches timestamps apply only to IPv4 traffic.
-

To configure RPM using the J-Web interface:

1. Select **Troubleshoot > RPM > Configure RPM**.
2. In the **Configure RPM** page, enter information as specified in [Table 7 on page 87](#).
  - a. Click **Add** to set up the **Owner Name** and **Performance Probe Tests**.
  - b. Select a probe owner from **Probe Owners** list and click **Delete** to remove the selected probe owner
  - c. Double-click one of the probe owners in **Probe Owners** list to display the list of performance probe tests.
  - d. Double-click one of the performance probe tests to edit the test parameters.
3. Enter the **Maximum Number of Concurrent Probes** and specify the **Probe Servers**.
4. Click **Apply** to apply the RPM probe settings.

**Table 7: RPM Probe Owner, Concurrent Probes, and Probe Servers Configuration Fields**

| Field                               | Function                                                                                                                                                             | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Probe Owners                        | Identifies a owner for whom 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. | <ol style="list-style-type: none"> <li>1. Click <b>Add</b> and type an owner name.</li> <li>2. In <b>Performance Probe Tests</b>, click <b>Add</b> to define the RPM test parameters. See <a href="#">Table 8 on page 88</a> for information on configuring RPM test parameters.</li> <li>3. Click <b>OK</b> to save the settings or <b>Cancel</b> to exit from the window without saving the changes.</li> </ol>                                                              |
| Maximum Number of Concurrent Probes | Specifies the maximum number of concurrent probes allowed.                                                                                                           | Type a number from 1 through 500.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Probe Servers                       | Specifies the servers that act as receivers and transmitters for the probes.                                                                                         | Set up the following servers: <ul style="list-style-type: none"> <li>• TCP Probe Server—Specifies the port on which the device is to receive and transmit TCP probes. Type the number 7 (a standard TCP port number) or a port number from 49160 through 65535.</li> <li>• UDP Probe Server—Specifies the port on which the device is to receive and transmit UDP probes. Type the number 7 (a standard TCP port number) or a port number from 49160 through 65535.</li> </ul> |

Table 8: Performance Probe Tests Configuration Fields

| Field                      | Function                                                                                                                                         | Your Action                                                                                                                                                                                                                                                                                             |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Identification</b>      |                                                                                                                                                  |                                                                                                                                                                                                                                                                                                         |
| Test Name                  | Identifies the RPM test.                                                                                                                         | Type a test name.                                                                                                                                                                                                                                                                                       |
| Target (Address or URL)    | Specifies the IP address or the URL of the probe target.                                                                                         | Type the IP address in dotted decimal notation or the URL of the probe target. If the target is a URL, type a fully formed URL that includes <b>http://</b> .                                                                                                                                           |
| Source Address             | Specifies the IP address to be used as the probe source address.                                                                                 | Type the source address to be used for the probe. If you do not supply this value, the packet uses the outgoing interface's address as the probe source address.                                                                                                                                        |
| Routing Instance           | Specifies the routing instance over which the probe is sent.                                                                                     | Type the routing instance name. The routing instance applies only to <b>icmp-ping</b> and <b>icmp-ping-timestamp</b> probe types. The default routing instance is <b>inet.0</b> .                                                                                                                       |
| History Size               | Specifies the number of probe results to be saved in the probe history.                                                                          | Type a number from 0 through 255. The default history size is 50.                                                                                                                                                                                                                                       |
| <b>Request Information</b> |                                                                                                                                                  |                                                                                                                                                                                                                                                                                                         |
| Probe Type                 | Specifies the type of probe to send as part of the test.                                                                                         | Select a probe type from the list: <ul style="list-style-type: none"> <li>• <b>http-get</b></li> <li>• <b>http-get-metadata</b></li> <li>• <b>icmp-ping</b></li> <li>• <b>icmp-ping-timestamp</b></li> <li>• <b>tcp-ping</b></li> <li>• <b>udp-ping</b></li> <li>• <b>udp-ping-timestamp</b></li> </ul> |
| Interval                   | Sets the wait time (in seconds) between probe transmissions.                                                                                     | Type a number from 1 through 255 .                                                                                                                                                                                                                                                                      |
| Test Interval              | Sets the wait time (in seconds) between tests.                                                                                                   | Type a number from 0 through 86400 .                                                                                                                                                                                                                                                                    |
| Probe Count                | Sets the total number of probes to be sent for each test.                                                                                        | Type a number from 1 through 15.                                                                                                                                                                                                                                                                        |
| Moving Average Size        | Specifies the number of samples to be used in the statistical calculation operations to be performed across a number of the most recent samples. | Type a number from 0 through 255.                                                                                                                                                                                                                                                                       |

Table 8: Performance Probe Tests Configuration Fields (*continued*)

| Field                           | Function                                                                                                                                                                                                                                                                                                           | Your Action                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Destination Port                | Specifies the TCP or UDP port to which probes are sent.<br><br>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 network devices configured to receive and transmit RPM probes on the same TCP or UDP port. | Type the number 7 (a standard TCP or UDP port number) or a port number from 49160 through 65535. |
| DSCP Bits                       | Specifies the Differentiated Services code point (DSCP) bits. This value must be a valid 6-bit pattern.                                                                                                                                                                                                            | Type a valid 6-bit pattern.                                                                      |
| Data Size                       | Specifies the size (in bytes) of the data portion of the ICMP probes.                                                                                                                                                                                                                                              | Type a number from 0 through 65507.                                                              |
| Data Fill                       | Specifies the hexadecimal value of the data portion of the ICMP probes.                                                                                                                                                                                                                                            | Type a hexadecimal value from 1h through 800h .                                                  |
| <b>Hardware Timestamp</b>       |                                                                                                                                                                                                                                                                                                                    |                                                                                                  |
| One Way Hardware Timestamp      | Enables one-way hardware timestamp.                                                                                                                                                                                                                                                                                | To enable timestamping, select the check box.                                                    |
| Destination Interface           | Enables hardware timestamp on the specified interface.                                                                                                                                                                                                                                                             | Select an interface from the list.                                                               |
| <b>Maximum Probe Thresholds</b> |                                                                                                                                                                                                                                                                                                                    |                                                                                                  |
| Successive Lost Probes          | Sets the number of probes that can be lost successively, if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                                                 | Type a number from 0 through 15.                                                                 |
| Lost Probes                     | Sets the number of probes that can be lost , if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                                                             | Type a number from 0 through 15.                                                                 |
| Round Trip Time                 | Sets the round-trip time (in microseconds), from the switch to the remote server, if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                        | Type a number from 0 through 60000000.                                                           |
| Jitter                          | Sets the jitter (in microseconds), if exceeded, triggers a probe failure and generates a system log message.                                                                                                                                                                                                       | Type a number from 0 through 60000000.                                                           |

Table 8: Performance Probe Tests Configuration Fields (*continued*)

| Field                              | Function                                                                                                                                                     | Your Action                                                                                                                                                           |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Standard Deviation                 | Sets the maximum allowable standard deviation (in microseconds), if exceeded, triggers a probe failure and generates a system log message.                   | Type a number from 0 through 60000000.                                                                                                                                |
| Egress Time                        | Sets the one-way time (in microseconds), from the switch to the remote server, if exceeded, triggers a probe failure and generates a system log message.     | Type a number from 0 through 60000000.                                                                                                                                |
| Ingress Time                       | Sets the one-way time (in microseconds), from the remote server to the switch, if exceeded, triggers a probe failure and generates a system log message.     | Type a number from 0 through 60000000 (microseconds).                                                                                                                 |
| Jitter Egress Time                 | Sets the outbound-time jitter (in microseconds), if exceeded triggers a probe failure and generates a system log message.                                    | Type a number from 0 through 60000000.                                                                                                                                |
| Jitter Ingress Time                | Sets the inbound-time jitter (in microseconds), if exceeded, triggers a probe failure and generates a system log message.                                    | Type a number from 0 and 60000000.                                                                                                                                    |
| Egress Standard Deviation          | Sets the maximum allowable standard deviation of outbound times (in microseconds), if exceeded, triggers a probe failure and generates a system log message. | Type a number from 0 through 60000000.                                                                                                                                |
| Ingress Standard Deviation         | Sets the maximum allowable standard deviation of inbound times (in microseconds), if exceeded, triggers a probe failure and generates a system log message.  | Type a number from 0 through 60000000.                                                                                                                                |
| <b>Traps</b>                       |                                                                                                                                                              |                                                                                                                                                                       |
| Egress Jitter Exceeded             | Generates SNMP traps when the threshold for jitter in outbound time is exceeded.                                                                             | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Egress Standard Deviation Exceeded | Generates SNMP traps when the threshold for standard deviation in outbound times is exceeded.                                                                | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |



Table 8: Performance Probe Tests Configuration Fields (*continued*)

| Field                               | Function                                                                                        | Your Action                                                                                                                                                           |
|-------------------------------------|-------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Egress Time Exceeded                | Generates SNMP traps when the threshold for maximum outbound time is exceeded.                  | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Ingress Jitter Exceeded             | Generates SNMP traps when the threshold for jitter in inbound time is exceeded.                 | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Ingress Standard Deviation Exceeded | Generates SNMP traps when the threshold for standard deviation in inbound times is exceeded.    | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Ingress Time Exceeded               | Generates SNMP traps when the threshold for maximum inbound time is exceeded.                   | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Jitter Exceeded                     | Generates SNMP traps when the threshold for jitter in round-trip time is exceeded.              | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Probe Failure                       | Generates SNMP traps when the threshold for the number of successive lost probes is exceeded.   | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| RTT Exceeded                        | Generates SNMP traps when the threshold for maximum round-trip time is exceeded.                | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Standard Deviation Exceeded         | Generates SNMP traps when the threshold for standard deviation in round-trip times is exceeded. | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Test Completion                     | Generates SNMP traps when a test is completed.                                                  | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |
| Test Failure                        | Generates SNMP traps when the threshold for the total number of lost probes is exceeded.        | <ul style="list-style-type: none"> <li>To enable SNMP traps for this condition, select the check box.</li> <li>To disable SNMP traps, clear the check box.</li> </ul> |

- Related Documentation**
- [Configuring SNMP \(J-Web Procedure\) on page 75](#)
  - [Viewing Real-Time Performance Monitoring Information on page 302](#)

## Configuring the Interface for RPM Timestamping for Client/Server on an EX Series Switch (CLI Procedure)

---

Use real-time performance monitoring (RPM) to configure active probes to track and monitor traffic across the network and to investigate network problems. To configure basic RPM probes on the EX Series switch, you must configure the probe owner, the test, and the specific parameters of the RPM probe.

You can also set a timestamp to improve the measurement of latency or jitter. The probe is timestamped by the device originating the probe (the RPM client). If you do not enable hardware timestamps, the timer values are set. You should configure both the RPM client (the requester) and the RPM server (the responder) to timestamp the RPM packets. However, if the RPM server does not support hardware timestamps, RPM can only report the round-trip measurements.

Timestamps apply only to IPv4 traffic.

You can enable hardware timestamps for the following RPM probe types:

- **icmp-ping**
- **icmp-ping-timestamp**
- **udp-ping**
- **udp-ping-timestamp**

To configure RPM probes and enable hardware timestamping:

1. Specify the probe owner:

```
[edit services rpm]
user@switch# set probe owner
```

2. Specify a test name. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.

```
[edit services rpm probe owner]
user@switch# set test test-name
```

3. Specify the packet and protocol contents of the probe:

```
[edit services rpm probe owner test test-name]
user@switch# set probe-type type
```

4. Specify the destination IPv4 address to be used for the probes:

```
[edit services rpm probe owner test test-name]
user@switch# set target address
```

5. Specify the number of probes within a test:

```
[edit services rpm probe owner test test-name]
user@switch# set probe-count count
```

6. Specify the time, in seconds, to wait between sending packets:

```
[edit services rpm probe owner test test-name]
```

- ```
user@switch# set probe-interval interval
```
7. Specify the time, in seconds, to wait between tests:


```
[edit services rpm probe owner test test-name]
user@switch# set test-interval interval
```
 8. Specify the source IP address to be used for probes. If the source IP address is not one of the switch's assigned addresses, the packet uses the outgoing interface's address as its source.


```
[edit services rpm probe owner test test-name]
user@switch# set source-address address
```
 9. Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.


```
[edit services rpm probe owner test test-name]
user@switch# set dscp-code-point dscp-bits
```
 10. If you are using ICMP probes, specify the size of the data portion of ICMP probes:


```
[edit services rpm probe owner test test-name]
user@switch# set data-size size
```
 11. Enable hardware timestamping of RPM probe messages:


```
[edit services rpm probe owner test test-name]
user@switch# set hardware-timestamp
```

**Related
Documentation**

- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 85](#)
- [Understanding Real-Time Performance Monitoring on EX Series Switches on page 3](#)

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Ethernet interfaces on Juniper Networks EX Series Ethernet Switches and Juniper Networks Junos OS for EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM).

This topic describes these tasks:

1. [Creating the Maintenance Domain on page 93](#)
2. [Configuring the Maintenance Domain MIP Half Function on page 94](#)
3. [Creating a Maintenance Association on page 94](#)
4. [Configuring the Continuity Check Protocol on page 95](#)
5. [Configuring a Maintenance Association End Point on page 95](#)
6. [Configuring a Connectivity Fault Management Action Profile on page 96](#)
7. [Configuring the Linktrace Protocol on page 96](#)

Creating the Maintenance Domain

A maintenance domain comprises network entities such as operators, providers, and customers. To enable connectivity fault management (CFM) on an Ethernet interface, you must create a maintenance domains, maintenance associations, and MEPS.

To create a maintenance domain:

1. Specify a name for the maintenance domain:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set maintenance-domain domain-name
```

2. Specify a format for the maintenance domain name. If you specify **none**, no name is configured:

- A plain ASCII character string
- A domain name service (DNS) format
- A media access control (MAC) address plus a two-octet identifier in the range 0 through 65,535
- **none**

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set name-format format
```

For example, to specify the name format as MAC address plus a two-octet identifier:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set name-format mac+2oct
```

3. Configure the maintenance domain level, which is used to indicate the nesting relationship between this domain and other domains. Use a value from 0 through 7:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set level level
```

Configuring the Maintenance Domain MIP Half Function

MIP Half Function (MHF) divides the maintenance association intermediate point (MIP) functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loop-back and link-trace messages to help isolate faults. Whenever a MIP is configured, the MIP half function value for all maintenance domains and maintenance associations must be the same.

To configure the MIP half function:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set mip-half-function (none | default | explicit)
```

Creating a Maintenance Association

In a CFM maintenance domain, each service instance is called a maintenance association.

To create a maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set maintenance-association ma-name
```

Configuring the Continuity Check Protocol

The continuity check protocol is used for fault detection by a maintenance association end point (MEP) within a maintenance association. The MEP periodically sends continuity check multicast messages. The receiving MEPs use the continuity check messages (CCMs) to build a MEP database of all MEPs in the maintenance association.

To configure the continuity check protocol:

1. Enable the continuity check protocol:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name]
user@switch# set continuity-check
```

2. Specify the continuity check hold interval. The hold interval is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name continuity-check]
user@switch# set hold-interval number
```

3. Specify the CCM interval. The interval is the time between the transmission of CCMs. You can specify 10 minutes (10m), 1 minute (1m), 10 seconds (10s), 1 second (1s), 100 milliseconds (100ms), or 10 milliseconds (10ms).

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name continuity-check]
user@switch# set interval number
```

4. Specify the number of CCMs (that is, protocol data units) that can be lost before the MEP is marked as down. The default number of protocol data units (PDUs) is 3.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name continuity-check]
user@switch# set loss-threshold number
```

Configuring a Maintenance Association End Point

To configure a maintenance association end point:

1. Specify an ID for the MEP. The value can be from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name]
user@switch# set mep mep-id
```

2. Enable maintenance endpoint automatic discovery if you want to have the MEP accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name mep mep-id]
user@switch# set auto-discovery
```

3. You can specify that CFM packets (CCMs) be transmitted only in one direction for the MEP, that is, the direction be set as **down** so that CCMs are transmitted only out of (not into) the interface configured on this MEP.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name mep mep-id]
user@switch# set direction down
```

- Specify the logical interface to which the MEP is attached. It can be either an access interface or a trunk interface. If you specify a trunk interface, the VLAN associated with that interface must have a VLAN ID.



NOTE: You cannot associate an access interface that belongs to multiple VLANs with the MEP.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name mep mep-id]
user@switch# set interface interface-name
```

- You can configure a remote MEP from which CCMs are expected. If autodiscovery is not enabled, the remote MEP must be configured under the **mep** statement. If the remote MEP is not configured under the **mep** statement, the CCMs from the remote MEP are treated as errors.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name maintenance-association ma-name mep mep-id]
user@switch# set remote-mep mep-id
```

Configuring a Connectivity Fault Management Action Profile

You can configure an action profile and specify the action to be taken when any of the configured events occur. Alternatively, you can configure an action profile and specify default actions when connectivity to a remote MEP fails.

To configure an action profile:

- Specify a name for an action profile:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set action-profile profile-name
```

- Configure the action of the action profile:

```
[edit protocols oam ethernet connectivity-fault-management action-profile
profile-name]
user@switch# set action interface-down
```

- Configure one or more events under the action profile, the occurrence of which will trigger the corresponding action to be taken:

```
[edit protocols oam ethernet connectivity-fault-management action-profile
profile-name]
user@switch# set event event
```

See [Junos OS Network Interfaces Configuration Guide](#)

Configuring the Linktrace Protocol

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between a MEP and a MIP under the same maintenance domain.

To configure the linktrace protocol:

1. Configure the linktrace path age timer. If no response to a linktrace request is received, the request and response entries are deleted after the age timer expires:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set linktrace age time
```

2. Configure the number of linktrace reply entries to be stored per linktrace request:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set linktrace path-database-size path-database-size
```

Related Documentation

- [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 68](#)
- [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch on page 22](#)
- [Junos OS Network Interfaces Configuration Guide](#)

Configuring Ethernet OAM Link Fault Management (CLI Procedure)

Ethernet OAM link fault management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across point-to-point Ethernet links either directly or through repeaters.

To configure Ethernet OAM LFM using the CLI:

1. Enable IEEE 802.3ah OAM support on an interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name
```



NOTE: You can configure Ethernet OAM LFM on aggregated interfaces.



NOTE: The remaining steps are optional. You can choose which of these features to configure for Ethernet OAM LFM on your switch.

2. Specify whether the interface or the peer initiates the discovery process by configuring the link discovery mode to **active** or **passive** (**active** = interface initiates; **passive** = peer initiates):

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name link-discovery active
```

3. Configure a periodic OAM PDU-sending interval (in milliseconds) for fault detection:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface pdu-interval interval
```

4. Specify the number of OAM PDUs that an interface can miss before the link between peers is considered down:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name pdu-threshold threshold-value
```

5. Configure event threshold values on an interface for the local errors that trigger the sending of link event TLVs:

- Set the threshold value (in seconds) for sending frame-error events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-error count
```

- Set the threshold value (in seconds) for sending frame-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period count
```

- Set the threshold value (in seconds) for sending frame-period-summary events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds frame-period-summary count
```

- Set the threshold value (in seconds) for sending symbol-period events or taking the action specified in the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name event-thresholds symbol-period count
```



NOTE: You can disable the sending of link event TLVs.

To disable the sending of link event TLVs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name negotiation-options no-allow-link-events
```

6. Create an action profile to define event fault flags and thresholds to be taken when the link fault event occurs. Then apply the action profile to one or more interfaces. (You can also apply multiple action profiles to a single interface.)

- a. Name the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name
```

- b. Specify actions to be taken by the system when the link fault event occurs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name action syslog
user@switch# set action-profile profile-name action link-down
```

- c. Specify events for the action profile:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set action-profile profile-name event link-adjacency-loss
```



NOTE: For each action profile, you must specify at least one link event and one action. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all actions are executed. You can set a low threshold for a specific action such as logging the error and set a high threshold for another action such as system logging.

7. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Set the remote DTE in loopback

mode (the remote DTE must support remote-loopback mode) and then enable remote loopback support for the local interface.

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name remote-loopback
user@switch# set interface interface-name negotiation-options allow-remote-loopback
```

**Related
Documentation**

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72](#)
- [Understanding Ethernet OAM Link Fault Management on page 21](#)

Configuring Interfaces for Uplink Failure Detection (CLI Procedure)

You can configure uplink failure detection on EX Series switches to help ensure balanced traffic flow. Using this feature, switches can monitor and detect link failure on uplink interfaces and can propagate the failure to downlink interfaces so that servers connected to those downlink interfaces can switch over to secondary interfaces.

Follow these configuration guidelines:

- You can configure a maximum of 48 groups for each switch.
- You can configure a maximum of 48 uplink interfaces and 48 downlink interfaces in each group.
- You can configure physical links and logical links in separate groups.
- Ensure that all the interfaces in the group are up. If the interfaces are down, uplink failure detection does not work.



NOTE: Routed VLAN interfaces (RVIs) cannot be configured as uplink interfaces to be monitored.

To configure uplink failure detection on a switch:

1. Specify a name for the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name
```

2. Add an uplink interface to the group:

```
[edit protocols]
user@switch# set uplink-failure-detection group group-name link-to-monitor interface-name
```

3. Repeat Step 2 for adding each uplink interface to the group.



NOTE: An interface can be configured as link-to-monitor in multiple groups.

4. Add a downlink interface to the group:

```
[edit protocols]
```

```
user@switch# set uplink-failure-detection group group-name link-to-disable interface-name
```

5. Repeat Step 4 for adding each downlink interface to the group.



NOTE: After you have configured a group, use the `show uplink-failure-detection group group-name` command to verify that all interfaces in the group are up.

**Related
Documentation**

- [Verifying That Uplink Failure Detection Is Working Correctly on page 302](#)
- [Understanding Uplink Failure Detection on page 27](#)

Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure)

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging service-level agreements (SLAs). By default, Ethernet frame delay measurement uses software for timestamping and delay calculations. You can configure an EX Series switch to perform and display Ethernet frame delay measurements on Ethernet interfaces. The switches support software-assisted timestamping.

Before you can begin configuring MEP interfaces to support Ethernet frame delay measurements on switches, ensure that you have:

- Configured Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) correctly
- Enabled distributed periodic packet management (PPM) (distributed PPM is enabled by default)

To configure MEP interfaces on switches to support Ethernet frame delay measurements:

1. Enable the Ethernet frame delay measurement by issuing the **monitor ethernet delay-measurement** operational mode command. In this command, you must specify one measurement type (either one-way or two-way measurement), and you must specify either the unicast MAC address of the peer MEP or its numeric identifier.

Optionally, you can also specify the following parameters:

- Number of frames to send to the peer MEP (**count count**)
- Number of seconds to wait between sending frames (**wait time**)
- Priority value of the delay measurement request frame (**priority value**)
- Size of the data in the data TLV of the request packet (**size value**)
- Suppression of the insertion of the session ID TLV in the request packet (**no-session-id-tlv**)

```
user@switch> monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name one-way mep remote-mep-id count count wait
time priority value size value no-session-id-tlv
```

Related Documentation

- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 103](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 24](#)

Configuring One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging service-level agreements (SLAs). You can configure the frame delay measurements in either a one-way mode or a two-way (round-trip) mode to gather frame delay statistics. For one-way Ethernet frame delay measurement, clocks at the local and remote MEPs need to be synchronized. However, clock synchronization is not required for two-way Ethernet frame delay measurement.

Before you begin configuring one-way Ethernet frame delay measurements on two EX Series switches, ensure that you have:

- Configured Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) correctly on both the switches
- Synchronized the system clocks of both the switches

To configure one-way Ethernet frame delay measurements:

1. Configure the maintenance domain, maintenance association, and MEP ID on both the switches.

2. From either switch, start a one-way Ethernet frame delay measurement:

```
user@switch> monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name one-way mep remote-mep-id count count wait
time
```

You can view the result on the other switch:

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md-name maintenance-association ma-name local-mep mep-id
remote-mep mep-id
```

Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 103](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 24](#)

Configuring an Iterator Profile on a Switch (CLI Procedure)

Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor service-level agreements (SLAs). You can create an iterator profile with its parameters to periodically transmit SLA measurement packets in the form of ITU-Y.1731-compliant frames for two-way delay measurement.

To create an iterator profile:

1. Specify a name for an SLA iterator profile—for example, *i1*:

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
user@switch# edit sla-iterator-profiles i1
```

2. (Optional) Configure the cycle time, which is the time (in milliseconds) between back-to-back transmissions of SLA frames.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set cycle-time cycle-time-value
```

3. (Optional) Configure the iteration period, which indicates the maximum number of cycles per iteration (the number of connections registered to an iterator cannot exceed this value).

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set iteration-period iteration-period-value
```

4. Configure the measurement type as two-way delay measurement.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
user@switch# set measurement-type two-way-delay
```

5. (Optional) Configure the calculation weight for delay.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
```

```
user@switch# set calculation-weight delay delay-value
```

6. (Optional) Configure the calculation weight for delay variation.

```
[edit protocols oam ethernet connectivity-fault-management performance-monitoring
sla-iterator-profiles i1]
```

```
user@switch# set calculation-weight delay-variation delay-variation-value
```

7. Configure a remote MEP with the iterator profile.

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
md-name maintenance-association ma-name mep mep-id remote-mep
remote-mep-id]
```

```
user@switch# set sla-iterator-profiles i1
```

Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 103](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 24](#)

Triggering an Ethernet Frame Delay Measurement Session on a Switch

To trigger Ethernet frame delay measurement, use the **monitor ethernet delay-measurement** operational command and specify the following values:

- Either one-way (**one-way**) or two-way (**two-way**) measurement
- Either the MAC address (**remote-mac-address**) or the MEP ID (**mep**) of the remote host
- The maintenance domain (**maintenance-domain**)
- The maintenance association (**maintenance-association**)
- (Optional) Any or all of these options: **count**, **size**, **wait**, **no-session-id-tlv**, **priority**

For example:

```
user@switch> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10 size 50 wait 5 no-session-id-tlv
priority 1
```

Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)

- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 24](#)

Configuring Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure)

Ethernet frame delay measurement is a useful tool for providing performance statistics or supporting or challenging service-level agreements (SLAs). You can configure the frame delay measurements in either a one-way mode or a two-way (round-trip) mode to gather frame delay statistics. For one-way Ethernet frame delay measurement, clocks at the local and remote MEPs need to be synchronized. However, clock synchronization is not required for two-way Ethernet frame delay measurement.

Before you begin configuring two-way Ethernet frame delay measurements on two EX Series switches, ensure that you have:

- Configured Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM) correctly on both the switches

To configure two-way Ethernet frame delay measurements:

1. Configure the maintenance domain, maintenance association, and MEP ID on both the switches.
2. From either switch, start a two-way Ethernet frame delay measurement:

```
user@switch> monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name two-way mep remote-mep-id count count wait
time
```

You can view the result on the other switch:

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md-name maintenance-association ma-name local-mep mep-id
remote-mep mep-id
```

Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)
- [Triggering an Ethernet Frame Delay Measurement Session on a Switch on page 103](#)
- [Understanding Ethernet Frame Delay Measurements on Switches on page 24](#)

CHAPTER 8

Configuration Statements: SNMP

- [\[edit snmp\] Configuration Statement Hierarchy on EX Series Switches on page 107](#)
- [address \(SNMP\) on page 114](#)
- [address-mask on page 115](#)
- [agent-address on page 115](#)
- [alarm \(SNMP RMON\) on page 116](#)
- [authorization on page 117](#)
- [bucket-size on page 117](#)
- [categories on page 118](#)
- [client-list on page 118](#)
- [client-list-name on page 119](#)
- [clients on page 119](#)
- [commit-delay on page 120](#)
- [community \(SNMP\) on page 121](#)
- [community \(SNMP RMON\) on page 122](#)
- [community-name \(SNMP\) on page 123](#)
- [contact \(SNMP\) on page 124](#)
- [description \(SNMP\) on page 125](#)
- [description \(SNMP RMON\) on page 126](#)
- [destination-port on page 126](#)
- [engine-id \(SNMP\) on page 127](#)
- [event \(SNMP\) on page 128](#)
- [falling-event-index on page 128](#)
- [falling-threshold on page 129](#)
- [falling-threshold on page 130](#)
- [falling-threshold-interval on page 131](#)
- [filter-duplicates on page 131](#)
- [filter-interfaces on page 132](#)
- [group \(Configuring Access Privileges\) on page 133](#)

- [group \(Associating a Security Name\) on page 134](#)
- [health-monitor on page 134](#)
- [history on page 135](#)
- [interface \(SNMP RMON History\) on page 136](#)
- [interface \(SNMP\) on page 136](#)
- [interval on page 137](#)
- [interval \(SNMP Health Monitor\) on page 137](#)
- [interval \(SNMP RMON\) on page 138](#)
- [location \(SNMP\) on page 138](#)
- [logical-system \(SNMP\) on page 139](#)
- [message-processing-model on page 140](#)
- [name on page 140](#)
- [nonvolatile on page 141](#)
- [notify on page 142](#)
- [notify-filter \(Applying to the Management Target\) on page 143](#)
- [notify-filter \(Configuring the Profile Name\) on page 143](#)
- [notify-view on page 144](#)
- [oid \(SNMP\) on page 144](#)
- [oid on page 145](#)
- [owner on page 145](#)
- [parameters on page 146](#)
- [port \(SNMP\) on page 146](#)
- [read-view on page 147](#)
- [request-type on page 148](#)
- [rising-event-index on page 149](#)
- [rising-threshold \(SNMP Health Monitor\) on page 150](#)
- [rising-threshold \(SNMP RMON\) on page 151](#)
- [rmon on page 151](#)
- [rmon on page 152](#)
- [routing-instance \(SNMP\) on page 153](#)
- [routing-instance \(SNMPv3\) on page 154](#)
- [sample-type on page 154](#)
- [security-level \(Defining Access Privileges\) on page 155](#)
- [security-level \(Generating SNMP Notifications\) on page 156](#)
- [security-model \(Access Privileges\) on page 157](#)
- [security-model \(Group\) on page 158](#)
- [security-model \(SNMP Notifications\) on page 159](#)

- [security-name \(Security Group\) on page 160](#)
- [security-name \(Community String\) on page 161](#)
- [security-name \(SNMP Notifications\) on page 162](#)
- [security-to-group on page 163](#)
- [snmp on page 163](#)
- [snmp on page 164](#)
- [snmp-community on page 165](#)
- [source-address \(SNMP\) on page 166](#)
- [startup-alarm on page 167](#)
- [syslog-subtag on page 167](#)
- [tag \(SNMPv3\) on page 168](#)
- [tag-list on page 168](#)
- [target-address on page 169](#)
- [target-parameters on page 170](#)
- [targets on page 171](#)
- [traceoptions \(SNMP\) on page 172](#)
- [trap-group on page 174](#)
- [trap-options on page 175](#)
- [type \(SNMP RMON\) on page 176](#)
- [type \(SNMPv3\) on page 177](#)
- [v3 on page 178](#)
- [vacm on page 180](#)
- [variable on page 181](#)
- [version \(SNMP\) on page 181](#)
- [view \(Associating a MIB View with a Community\) on page 182](#)
- [view \(Configuring a MIB View\) on page 183](#)
- [write-view on page 184](#)

[edit snmp] Configuration Statement Hierarchy on EX Series Switches

This topic lists supported and unsupported configuration statements in the **[edit snmp]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see *EX Series Switch Software Features Overview*

This topic lists:

- [Supported Statements in the \[edit snmp\] Hierarchy Level on page 108](#)
- [Unsupported Statements in the \[edit snmp\] Hierarchy Level on page 112](#)

Supported Statements in the [edit snmp] Hierarchy Level

The following hierarchy shows the **[edit snmp]** configuration statements supported on EX Series switches:

```
snmp {
  client-list list-name {
    address {
      restrict;
    }
  }
  community community-name {
    authorization (read-only | read-write);
    client-list-name list-name;
    clients {
      address <restrict>;
    }
    routing-instance instance-name;
    routing-instance instance-name {
      client-list-name list-name;
      clients {
        address <restrict>;
      }
    }
    view view-name;
  }
  contact contact-information;
  description description;
  engine-id {
    (local engine-id | use-default-ip-address | use-mac-address);
  }
  filter-duplicates;
  filter-interfaces {
    interfaces
    all-internal-interfaces;
    interface 1;
    interface 2;
  }
  health-monitor {
    falling-threshold percentage;
    idp {
      falling-threshold;
      interval seconds;
      rising-threshold;
    }
    interval seconds;
    rising-threshold percentage;
  }
  interface [ interface-names ];
  location location;
```

```

name system-name;
nonvolatile {
    commit-delay seconds;
}
rmon {
    alarm index {
        description description;
        falling-event-index index;
        falling-threshold integer;
        falling-threshold-interval seconds;
        interval seconds;
        request-type (get-next-request | get-request | walk-request);
        rising-event-index index;
        rising-threshold integer;
        sample-type (absolute-value | delta-value);
        startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
        syslog-subtag text-string;
        variable oid-variable;
    }
    event index {
        community community-name;
        description description;
        type (log | log-and-trap | none | snmptrap);
    }
    history index {
        bucket-size number;
        interface interface-name;
        interval seconds;
        owner owner-name;
    }
}
routing-instance-access {
    access-list {
        routing-instance-name <restrict>;
    }
}
traceoptions {
    file <files number> <match regular-expression> <size maximum-file-size>
        <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
trap-group group-name {
    categories {
        authentication;
        chassis;
        configuration;
        link;
        otn-alarms {
            alarm-name;
        }
        remote-operations;
        rmon-alarm;
        routing;
        services;
        sonet-alarms {

```

```

        alarm-name;
    }
    startup;
    vrrp-events;
}
destination-port port-number;
routing-instance instance-name;
routing-instance instance-name;
targets {
    address;
}
version (all | v1 | v2);
}
trap-options {
    agent-address outgoing-interface;
    enterprise-oid;
    routing-instance instance-name;
    routing-instance instance-name {
        source-address (address | lo0);
    }
    source-address address;
}
v3 {
    ... the v3 subhierarchy appears after the main [edit snmp] hierarchy level ...
}
view view-name {
    oid object-identifier <exclude | include>;
}
}

snmp {
    v3 {
        notify name {
            tag tag-name;
            type (inform | trap);
        }
        notify-filter profile-name {
            oid oid <exclude | include>;
        }
        snmp-community community-index {
            community-name community-name;
            context context-name;
            security-name security-name;
            tag tag-name;
        }
        target-address target-address-name {
            address address;
            address-mask address-mask;
            routing-instance routing-instance-name;
            port port-number;
            retry-count number;
            routing-instance routing-instance-name;
            tag-list tag-list;
            target-parameters parameter-name;
            timeout seconds;
        }
    }
}

```

```

target-parameters parameter-name {
  notify-filter profile-name;
  parameters {
    message-processing-model (v1 | v2c | v3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
usm {
  local-engine {
    user username {
      authentication-md5 {
        authentication-key password;
        authentication-password password;
      }
      authentication-none;
      authentication-sha {
        authentication-key password;
        authentication-password password;
      }
      privacy-3des {
        privacy-password password;
      }
      privacy-aes128 {
        privacy-password password;
      }
      privacy-des {
        privacy-password password;
      }
      privacy-none;
    }
  }
  remote-engine engine-id {
    user username {
      authentication-md5 {
        authentication-key password;
        authentication-password password;
      }
      authentication-none;
      authentication-sha {
        authentication-key
        authentication-password password;
      }
      privacy-3des {
        privacy-password password;
      }
      privacy-aes128 {
        privacy-password password;
      }
      privacy-des {
        privacy-password password;
      }
      privacy-none;
    }
  }
}

```

```

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

```

Unsupported Statements in the [edit snmp] Hierarchy Level

All statements in the **[edit snmp]** hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented with the following exceptions:

Table 9: Unsupported [edit snmp] Configuration Statements on EX Series Switches

Statement	Hierarchy
-----------	-----------

NOTE: Variables, such as *community-name*, are not shown in the statements or hierarchies.

Table 9: Unsupported [edit snmp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
logical-system	[edit snmp community] [edit snmp trap-group]

Table 9: Unsupported [edit snmp] Configuration Statements on EX Series Switches (*continued*)

Statement	Hierarchy
	[edit snmp trap-options] [edit snmp v3 target-address]
logical-systems-trap-filter	[edit snmp]

- Related Documentation**
- [Configuring SNMP \(J-Web Procedure\) on page 75](#)
 - *Network Management Administration Guide for Routing Devices*

address (SNMP)

Syntax	address <i>address</i> ;
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the SNMP target address.
Options	address —IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Address</i>

address-mask

Syntax	<code>address-mask <i>address-mask</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 on the QFX Series.
Description	Define and verify the source addresses for a group of target addresses for SNMP traps and informs.
Options	<i>address-mask</i> —Define a range of addresses.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Address Mask</i>

agent-address

Syntax	<code>agent-address outgoing-interface;</code>
Hierarchy Level	<code>[edit snmp trap-options]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the agent address of all SNMPv1 traps generated by this router or switch. Currently, the only option is outgoing-interface , which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.
Options	outgoing-interface —Value of the agent address of all SNMPv1 traps generated by this router or switch. The outgoing-interface option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap. Default: disabled (the agent address is not specified in SNMPv1 traps).
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Agent Address for SNMP Traps</i>

alarm (SNMP RMON)

Syntax	<pre>alarm <i>index</i> { <i>description</i> <i>description</i>; <i>falling-event-index</i> <i>index</i>; <i>falling-threshold</i> <i>integer</i>; <i>falling-threshold-interval</i> <i>seconds</i>; <i>interval</i> <i>seconds</i>; <i>request-type</i> (get-next-request get-request walk-request); <i>rising-event-index</i> <i>index</i>; <i>rising-threshold</i> <i>integer</i>; <i>sample-type</i> (absolute-value delta-value); <i>startup-alarm</i> (falling-alarm rising-alarm rising-or-falling alarm); <i>syslog-subtag</i> <i>syslog-subtag</i>; <i>variable</i> <i>oid-variable</i>; }</pre>
Hierarchy Level	[edit snmp rmon]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure RMON alarm entries.
Options	<i>index</i> —Identifies this alarm entry as an integer. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an RMON Alarm Entry and Its Attributes</i>• event (SNMP) on page 128• <i>Configuring RMON Alarms and Events</i>• <i>RMON MIB Event, Alarm, Log, and History Control Tables</i>• <i>Monitoring RMON MIB Tables</i>• <i>Understanding RMON</i>

authorization

Syntax	<code>authorization <i>authorization</i>;</code>
Hierarchy Level	[edit <code>snmp community <i>community-name</i></code>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Set the access authorization for SNMP Get , GetBulk , GetNext , and Set requests.
Options	<p><i>authorization</i>—Access authorization level:</p> <ul style="list-style-type: none"> read-only—Enable Get, GetNext, and GetBulk requests. read-write—Enable all requests, including Set requests. You must configure a view to enable Set requests. <p>Default: <code>read-only</code></p>
Required Privilege Level	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring SNMP Communities

bucket-size

Syntax	<code>bucket-size <i>number</i>;</code>
Hierarchy Level	[edit <code>snmp rmon history</code>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.
Default	50
Options	<i>number</i> —Number of discrete samples of Ethernet statistics requested.
Required Privilege Level	<p><code>snmp</code>—To view this statement in the configuration.</p> <p><code>snmp-control</code>—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring SNMP (J-Web Procedure) on page 75 Junos OS Network Management Configuration Guide

categories

Syntax	<pre>categories { category; }</pre>
Hierarchy Level	[edit snmp trap-group <i>group-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	Define the types of traps that are sent to the targets of the named trap group.
Default	If you omit the categories statement, all trap types are included in trap notifications.
Options	category —Name of a trap type: authentication , chassis , configuration , link , remote-operations , rmon-alarm , routing , services , sonet-alarms , startup , or vrpp-events .
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Groups</i>

client-list

Syntax	<pre>client-list <i>client-list-name</i> { ip-addresses; }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for QFX Series switches.
Description	Define a list of SNMP clients.
Options	client-list-name —Name of the client list. ip-addresses —IP addresses of the SNMP clients to be added to the client list,
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Adding a Group of Clients to an SNMP Community</i>

client-list-name

Syntax	<code>client-list-name <i>client-list-name</i>;</code>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Add a client list or prefix list to an SNMP community.
Options	<i>client-list-name</i> —Name of the client list or prefix list.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Adding a Group of Clients to an SNMP Community</i>

clients

Syntax	<pre>clients { address <restrict>; }</pre>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.
Default	If you omit the clients statement, all SNMP clients using this community string are authorized to access the router.
Options	<p>address—Address of an SNMP client that is authorized to access this router. You must specify an address, not a hostname. To specify more than one client, include multiple address options.</p> <p>restrict—(Optional) Do not allow the specified SNMP client to access the router.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring SNMP Communities</i>

commit-delay

Syntax	commit-delay <i>seconds</i> ;
Hierarchy Level	[edit snmp nonvolatile]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the timer for the SNMP Set reply and start of the commit.
Options	seconds —Delay between an affirmative SNMP Set reply and start of the commit. Default: 5 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Commit Delay Timer</i>


community (SNMP)

Syntax	<pre>community <i>community-name</i> { authorization <i>authorization</i>; client-list-name <i>client-list-name</i>; clients { address restrict; } view <i>view-name</i>; }</pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.</p> <p>The SNMP client application specifies an SNMP community name in Get, GetBulk, GetNext, and Set SNMP requests.</p>
Default	If you omit the community statement, all SNMP requests are denied.
Options	<p>community-name—Community string. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring SNMP Communities</i>

community (SNMP RMON)

Syntax	<code>community <i>community-name</i>;</code>
Hierarchy Level	[edit snmp rmon event <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The trap group that is used when generating a trap (if eventType is configured to send traps). If that trap group has the rmon-alarm trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of eventCommunity). If nothing is configured, traps are sent to each group with the rmon-alarm category set.
Options	community-name —Identifies the trap group that is used when generating a trap if the event is configured to send traps.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an RMON Event Entry and Its Attributes</i>

community-name (SNMP)

Syntax	<code>community-name <i>community-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 snmp-community <i>community-index</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11. for the QFX Series.
Description	Define an SNMP community to authorize SNMPv1 or SNMPv2c clients in an SNMPv3 system. When you configure a community in SNMPv3, you can also specify a security name. The access privileges associated with the security name determine which MIB objects are available and which operations (read, write, or notify) are allowed on those objects.
Options	<i>community-name</i> —Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose the name in quotation marks (" ").
<div>  <p>NOTE: Community names must be unique. You cannot configure the same community name at the <code>[edit snmp community]</code> and <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy levels.</p> <p>The community name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level is encrypted and not displayed in the command-line interface (CLI).</p> </div>	
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring the SNMPv3 Community</i>

contact (SNMP)

Syntax	<code>contact <i>contact</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define the value of the MIB II sysContact object, which is the contact person for the managed system.
Options	contact —Name of the contact person. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the System Contact on a Device Running Junos OS</i>

description (SNMP)

Syntax	<code>description <i>description</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define the value of the MIB II sysDescription object, which is the description of the system being managed.
Default	By default, the sysDescription object includes the following information: Juniper Networks, Inc. <i>platform</i> , <i>build</i> , Build date: <i>date</i> UTC Copyright (c) <i>date-range</i> Juniper Networks, Inc. For example: sysDescr.0 = Juniper Networks, Inc. m7i internet router, kernel JUNOS 13.2-20130530_ib_13_3_psd.1, Build date: 2013-05-30 22:48:07 UTC Copyright (c) 1996-2013 Juniper Networks, Inc.
Options	<i>description</i> —System description. If the name includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the System Description on a Device Running Junos OS</i>


description (SNMP RMON)

Syntax	<code>description</code> <i>description</i> ;
Hierarchy Level	[edit snmp rmon alarm <i>index</i>], [edit snmp rmon event <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Text description of alarm or event.
Options	<i>description</i> —Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Description</i>• <i>Configuring an RMON Event Entry and Its Attributes</i>

destination-port

Syntax	<code>destination-port</code> <i>port-number</i> ;
Hierarchy Level	[edit snmp trap-group]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Assign a trap port number other than the default.
Default	If you omit this statement, the default port is 162.
Options	<i>port-number</i> —SNMP trap port number.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Groups</i>

engine-id (SNMP)

Syntax	engine-id { (local <i>engine-id-suffix</i> use-default-ip-address use-mac-address); }
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.1 for EX Series switches.
Description	The local engine ID is defined as the administratively unique identifier of an SNMPv3 engine, and is used for identification, not for addressing. There are two parts of an engine ID: prefix and suffix. The prefix is formatted according to the specifications defined in RFC 3411, <i>An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks</i> . You can configure the suffix here.
<div>  <p>NOTE: SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID. If you configure or change the engine ID, you must commit the new engine ID before you configure SNMPv3 users. Otherwise the keys generated from the configured passwords are based on the previous engine ID.</p> <p>For the engine ID, we recommend using the MAC address of the management port.</p> </div>	
Options	<p>local <i>engine-id-suffix</i>—Explicit setting for the engine ID suffix.</p> <p>use-default-ip-address—The engine ID suffix is generated from the default IP address.</p> <p>use-mac-address—The SNMP engine identifier is generated from the MAC address of the management interface on the router.</p> <p>Default: use-default-ip-address</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring the Local Engine ID

event (SNMP)

Syntax	<pre>event index { community community-name; description description; type type; }</pre>
Hierarchy Level	[edit snmp rmon]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure RMON event entries.
Options	<i>index</i> —Identifier for a specific event entry. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an RMON Event Entry and Its Attributes</i>• alarm on page 116

falling-event-index

Syntax	falling-event-index <i>index</i> ;
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The index of the event entry that is used when a falling threshold is crossed. If this value is zero, no event is triggered.
Options	<i>index</i> —Index of the event entry that is used when a falling threshold is crossed. Range: 0 through 65,535 Default: 0
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Falling Event Index or Rising Event Index</i>• rising-event-index on page 149

falling-threshold

Syntax	<code>falling-threshold <i>percentage</i>;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The lower threshold is expressed as a percentage of the maximum possible value for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the rising-threshold .
Options	<i>percentage</i> —The lower threshold for the alarm entry. Range: 1 through 100 Default: 70 percent of the maximum possible value
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Falling Threshold or Rising Threshold</i> • rising-threshold (SNMP Health Monitor) on page 150

falling-threshold

Syntax	<code>falling-threshold <i>integer</i>;</code>
Hierarchy Level	[edit snmp rmon alarm <i>index</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The lower threshold for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold, and the associated startup-alarm value is equal to falling-alarm value or rising-or-falling-alarm value. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the rising-threshold .
Options	integer —The lower threshold for the alarm entry. Range: -2,147,483,648 through 2,147,483,647 Default: 20 percent less than rising-threshold
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Falling Threshold or Rising Threshold</i>• rising-threshold (SNMP RMON) on page 151


falling-threshold-interval

Syntax	<code>falling-threshold-interval seconds;</code>
Hierarchy Level	<code>[edit snmp rmon alarm index]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Interval between samples when the rising threshold is crossed. Once the alarm crosses the falling threshold, the regular sampling interval is used.
Options	seconds —Time between samples, in seconds. Range: 1 through 2,147,483,647 seconds Default: 60 seconds
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Falling Threshold Interval</i> • interval (SNMP RMON) on page 138

filter-duplicates

Syntax	<code>filter-duplicates;</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Filter duplicate Get , GetNext , or GetBulk SNMP requests.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Filtering Duplicate SNMP Requests</i>

filter-interfaces

Syntax	<pre>filter-interfaces { all-internal-interfaces; interfaces <i>interface-names</i>{ interface 1; interface 2; } }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 9.4. Statement introduced in Junos OS Release 9.4 for EX Series Switches.
Description	Filter out information related to specific interfaces from the output of SNMP Get and GetNext requests performed on interface-related MIBs.
Options	<p>all-internal-interfaces—Filters out information from SNMP Get and GetNext requests for internal interfaces that include interfaces belonging to internal routing tables, such as <code>_juniper_private1_</code>, <code>_juniper_private2_</code>, <code>_juniper_private3_</code>, and <code>_juniper_private4_</code>. In addition, this option filters out information about the physical interface corresponding to a logical interface with a unit number from the internal routing table.</p> <p>interfaces—Specifies the interfaces to filter out from the output of SNMP Get and GetNext requests.</p>
	<div> NOTE: Starting with Release 12.1, Junos OS provides an except option (! operator) that enables you to filter out all interfaces except those interfaces that match the regular expressions prefixed with the ! mark.</div>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Filtering Interface Information Out of SNMP Get and GetNext Output</i>

group (Configuring Access Privileges)

```
Syntax  group group-name {
        (default-context-prefix | context-prefix context-prefix){
            security-model (any | usm | v1 | v2c) {
                security-level (authentication | none | privacy) {
                    notify-view view-name;
                    read-view view-name;
                    write-view view-name;
                }
            }
        }
    }
```

Hierarchy Level [edit snmp v3 vacm access]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Assign the security name to a group, and specify the SNMPv3 context applicable to the group. The **default-context-prefix** statement, when included, adds all the contexts configured on the device to the group, whereas the **context-prefix context-prefix** statement enables you to specify a context and to add that particular context to the group.

(Not applicable to the QFX Series.) When the context prefix is specified as default (for example, **context-prefix default**), the context associated with the master routing instance is added to the group. To specify a routing instance that is part of a logical system, specify it as **logical system/routing instance**. For example, to specify routing instance ri1 in logical system ls1, include **context-prefix ls1/ri1**.

The remaining statements under this hierarchy are explained separately.

Options *group-name*—SNMPv3 group name created for the SNMPv3 group.

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

Related Documentation

- *Configuring the Group*

group (Associating a Security Name)

Syntax	<code>group group-name;</code>
Hierarchy Level	<code>[edit snmp v3 vacm security-to-group security-model (usm v1 v2c) security-name security-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate a security name with a group composed of users with the same access privileges. The security name is used during authentication of SNMP messages, and is mapped to a username.
Options	<i>group-name</i> —Collection of SNMP security names that share the same SNMPv3 access privileges.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Group</i>

health-monitor

Syntax	<code>health-monitor { falling-threshold percentage; interval seconds; rising-threshold percentage; }</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure health monitoring. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Health Monitoring on Devices Running Junos OS</i>

history

Syntax	<pre>history <i>history-index</i> { bucket-size <i>number</i>; interface (SNMP RMON History) <i>interface-name</i>; interval <i>seconds</i>; owner <i>owner-name</i>; }</pre>
Hierarchy Level	[edit snmp rmon]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments. It collects statistics in accordance with user-configurable parameters.</p> <p>The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. The interface (SNMP RMON History) <i>interface-name</i> statement is mandatory. Other statements in the history group are optional.</p>
Default	Not configured.
Options	<p><i>history-index</i>—Identifies this history entry as an integer.</p> <p>Range: 1 through 655535</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP (J-Web Procedure) on page 75 • Junos OS Network Management Configuration Guide

interface (SNMP RMON History)

Syntax	<code>interface <i>interface-name</i>;</code>
Hierarchy Level	[edit snmp rmon history <i>history-index</i>]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Specify the interface to be monitored in the specified RMON history entry.</p> <p>Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created.</p>
Options	<i>interface-name</i> —Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP (J-Web Procedure) on page 75• Junos OS Network Management Configuration Guide

interface (SNMP)

Syntax	<code>interface [<i>interface-names</i>];</code>
Hierarchy Level	[edit <code>snmp</code>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the interfaces on which SNMP requests can be accepted.
Default	If you omit this statement, SNMP requests entering the router or switch through any interface are accepted.
Options	<i>interface-names</i> —Names of one or more logical interfaces.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring the Interfaces on Which SNMP Requests Can Be Accepted

interval

Syntax	<code>interval seconds;</code>
Hierarchy Level	[edit snmp rmon history]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the interval over which data is to be sampled for the specified interface.
Default	1800 sec
Options	<i>seconds</i> —Interval at which data is to be sampled for the specified interface.
Required Privilege	snmp—To view this statement in the configuration.
Level	snmp-control—To add this statement to the configuration.

interval (SNMP Health Monitor)

Syntax	<code>interval seconds;</code>
Hierarchy Level	[edit snmp health-monitor]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Interval between samples.
Options	<i>seconds</i> —Time between samples, in seconds. Range: 1 through 2147483647 seconds Default: 300 seconds
Required Privilege	snmp—To view this statement in the configuration.
Level	snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Interval</i>


interval (SNMP RMON)

Syntax	interval <i>seconds</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Interval between samples.
Options	seconds —Time between samples, in seconds. Range: 1 through 2,147,483,647 seconds Default: 60 seconds
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Interval</i>

location (SNMP)

Syntax	location <i>location</i> ;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Define the value of the MIB II sysLocation object, which is the physical location of the managed system.
Options	location —Location of the local system. You must enclose the name within quotation marks (" ").
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the System Location for a Device Running Junos OS</i>

logical-system (SNMP)

Syntax	<pre>logical-system <i>logical-system-name</i> { <i>routing-instance routing-instance-name</i>; }</pre>
Hierarchy Level	<pre>[edit snmp <i>community community-name</i>], [edit snmp <i>trap-group</i>], [edit snmp <i>trap-options</i>] [edit snmp <i>v3target-address target-address-name</i>]</pre>
Release Information	<p>Statement introduced in Junos OS Release 9.3</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The <code>logical-system</code> statement replaces the <code>logical-router</code> statement, and is backward-compatible with Junos OS Release 8.3 and later.</p> </div> </div>	
Description	<p>Specify a logical system name for SNMP v1 and v2c clients.</p> <p>Include at the <code>[edit snmp trap-options]</code> hierarchy level to specify a logical-system address as the source address of an SNMP trap.</p> <p>Include at the <code>[edit snmp v3 target-address]</code> hierarchy level to specify a logical-system name as the destination address for an SNMPv3 trap or inform.</p>
Options	<p><i>logical-system-name</i>—Name of the logical system.</p> <p><i>routing-instance routing-instance-name</i>—Statement to specify a routing instance associated with the logical system.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community</i> <i>Configuring the Trap Target Address</i>

message-processing-model

Syntax	message-processing-model (v1 v2c v3);
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameter-name</i> parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the message processing model to be used when generating SNMP notifications.
Options	v1—SNMPv1 message process model. v2c—SNMPv2c message process model. v3—SNMPv3 message process model.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Message Processing Model</i>

name

Syntax	name <i>name</i> ;
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the system name from the command-line interface.
Options	<i>name</i> —System name override.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the System Name</i>

nonvolatile

Syntax	<code>nonvolatile { commit-delay <i>seconds</i>; }</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. The commit-delay statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure options for SNMP Set requests. The statement is explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Commit Delay Timer</i>• commit-delay on page 120

notify

Syntax	<pre>notify <i>name</i> { tag <i>tag-name</i>; type (trap inform); }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. type inform option added in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Select management targets for SNMPv3 notifications as well as the type of notifications. Notifications can be either traps or informs.
Options	<p><i>name</i>—Name assigned to the notification.</p> <p><i>tag-name</i>—Notifications are sent to all targets configured with this tag.</p> <p><i>type</i>—Notification type is trap or inform. Traps are unconfirmed notifications. Informs are confirmed notifications.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Inform Notification Type and Target Address</i>• <i>Configuring the SNMPv3 Trap Notification</i>

notify-filter (Applying to the Management Target)

Syntax	<code>notify-filter <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-parameters <i>target-parameters-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the notify filter applied to a specific set of SNMPv3 target parameters. Target parameters are the message processing and security parameters for notifications sent to a target SNMP manager.
Options	<i>profile-name</i> —Name of the notify filter to apply to notifications.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Applying the Trap Notification Filter

notify-filter (Configuring the Profile Name)

Syntax	<code>notify-filter <i>profile-name</i> { oid <i>oid</i> (include exclude); }</code>
Hierarchy Level	<code>[edit snmp v3]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Specify a group of MIB objects for which you define access. The notify filter limits the type of traps or informs sent to the network management system.
Options	<p><i>profile-name</i>—Name assigned to the notify filter.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring the Trap Notification Filter • oid (SNMP) on page 144

notify-view

Syntax	<code>notify-view <i>view-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate the notify view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —Name of the view to which the SNMP user group has access.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring MIB Views</i>• <i>Configuring the Notify View</i>

oid (SNMP)

Syntax	<code>oid <i>oid</i> (include exclude);</code>
Hierarchy Level	<code>[edit snmp v3 notify-filter <i>profile-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects. This OID is a prefix that the represented MIB objects have in common.
Options	exclude —Exclude the subtree of MIB objects represented by the specified OID. include —Include the subtree of MIB objects represented by the specified OID. <i>oid</i> —Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Trap Notification Filter</i>

oid

Syntax	<code>oid <i>object-identifier</i> (exclude include);</code>
Hierarchy Level	[edit snmp view <i>view-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify an object identifier (OID) used to represent a subtree of MIB objects.
Options	<p>exclude—Exclude the subtree of MIB objects represented by the specified OID.</p> <p>include—Include the subtree of MIB objects represented by the specified OID.</p> <p><i>object-identifier</i>—OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring MIB Views

owner

Syntax	<code>owner <i>owner-name</i>;</code>
Hierarchy Level	[edit snmp rmon history]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the user or group responsible for this configuration.
Options	<p><i>owner-name</i>—The user or group responsible for this configuration.</p> <p>Range: 0 through 32 alphanumeric characters</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring SNMP (J-Web Procedure) on page 75 • Junos OS Network Management Configuration Guide

parameters

Syntax	<pre>parameters { message-processing-model (v1 v2c v3); security-level (none authentication privacy); security-model (usm v1 v2c); security-name security-name; }</pre>
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameters-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a set of target parameters for message processing and security. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Defining and Configuring the Trap Target Parameters</i>

port (SNMP)

Syntax	<pre>port port-number;</pre>
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure a UDP port number for an SNMP target.
Default	If you omit this statement, the default port is 162.
Options	<i>port-number</i> —Port number for the SNMP target.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Port</i>

read-view

Syntax	<code>read-view <i>view-name</i>;</code>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate the read-only view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —The name of the view to which the SNMP user group has access.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Read View</i>• <i>Configuring MIB Views</i>

request-type

Syntax	request-type (get-next-request get-request walk-request);
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Extend monitoring to a specific SNMP object instance (get-request), or extend monitoring to all object instances belonging to a MIB branch (walk-request), or extend monitoring to the next object instance after the instance specified in the configuration (get-next-request).
Options	get-next-request —Performs an SNMP get next request. get-request —Performs an SNMP get request. walk-request —Performs an SNMP walk request. Default: walk-request
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Request Type</i>• variable on page 181

rising-event-index

Syntax	rising-event-index <i>index</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Index of the event entry that is used when a rising threshold is crossed. If this value is zero, no event is triggered.
Options	<i>index</i> —Index of the event entry that is used when a rising threshold is crossed. Range: 0 through 65,535 Default: 0
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Falling Event Index or Rising Event Index</i>• falling-event-index on page 128

rising-threshold (SNMP Health Monitor)

Syntax	<code>rising-threshold <i>percentage</i>;</code>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 8.0. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The upper threshold is expressed as a percentage of the maximum possible value for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the falling-threshold .
Options	<i>percentage</i> —The lower threshold for the alarm entry. Range: 1 through 100 Default: 80 percent of the maximum possible value
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• falling-threshold on page 129• <i>Configuring the Falling Threshold or Rising Threshold</i>

rising-threshold (SNMP RMON)

Syntax	<code>rising-threshold <i>integer</i>;</code>
Hierarchy Level	<code>[edit snmp rmon alarm index]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Upper threshold for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold, and the associated startup alarm value is equal to the falling alarm or rising or falling alarm value. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the falling threshold.
Options	<i>integer</i> —The lower threshold for the alarm entry. Range: –2,147,483,648 through 2,147,483,647
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Falling Threshold or Rising Threshold</i> • falling-threshold on page 130

rmon

Syntax	<code>rmon { ... }</code>
Hierarchy Level	<code>[edit snmp]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure Remote Monitoring.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring an RMON Alarm Entry and Its Attributes</i>

rmon

Syntax	<pre>rmon { history <i>history-index</i> { interface (SNMP RMON History) <i>interface-name</i>; bucket-size <i>number</i>; interval <i>seconds</i>; owner <i>owner-name</i>; } }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>RMON is an existing feature of Junos OS.</p> <p>The RMON specification provides network administrators with comprehensive network fault diagnosis, planning, and performance tuning information. It delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.</p> <p>Junos OS supports RMON Statistics, History, Alarm, and Event groups. The EX Series documentation describes only the rmon history statement, which was added with this release.</p> <p>The statements are explained separately.</p>
Default	Disabled.
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring SNMP (J-Web Procedure) on page 75• Junos OS Network Management Configuration Guide

routing-instance (SNMP)

Syntax	<code>routing-instance <i>routing-instance-name</i>;</code>
Hierarchy Level	<code>[edit snmp community <i>community-name</i>],</code> <code>[edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>],</code> <code>[edit snmp trap-group <i>group</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Added to the <code>[edit snmp community <i>community-name</i>]</code> hierarchy level in Junos OS Release 8.4. Added to the <code>[edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>]</code> hierarchy level in Junos OS Release 9.1. Statement introduced in Junos OS Release 9.1 for EX Series switches.
Description	Specify a routing instance for SNMPv1 and SNMPv2 trap targets. All targets configured in the trap group use this routing instance. If the routing instance is defined within a logical system, include the logical-system <i>logical-system-name</i> statement at the <code>[edit snmp community <i>community-name</i>]</code> hierarchy level and specify the routing-instance statement under the <code>[edit snmp community <i>community-name</i> logical-system <i>logical-system-name</i>]</code> hierarchy level.
Options	<i>routing-instance-name</i> —Name of the routing instance.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring SNMP Trap Groups</i> • <i>Configuring the Source Address for SNMP Traps</i> • <i>Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community</i>

routing-instance (SNMPv3)

Syntax	<code>routing-instance <i>routing-instance-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-address <i>target-address-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 8.3. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify a routing instance for an SNMPv3 trap target.
Options	<p><i>routing-instance-name</i>—Name of the routing instance.</p> <p>To configure a routing instance within a logical system, specify the logical system name followed by the routing instance name. Use a slash (/) to separate the two names (for example, test-ls/test-ri). To configure the default routing instance on a logical system, specify the logical system name followed by default (for example, test-ls/default).</p>
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Trap Target Address</i>

sample-type

Syntax	<code>sample-type (absolute-value delta-value);</code>
Hierarchy Level	<code>[edit snmp rmon alarm <i>index</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Method of sampling the selected variable.
Options	<p>absolute-value—Actual value of the selected variable is used when comparing against the thresholds.</p> <p>delta-value—Difference between samples of the selected variable is used when comparing against the thresholds.</p>
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Sample Type</i>

security-level (Defining Access Privileges)

Syntax	<pre>security-level (authentication none privacy) { notify-view view-name; read-view view-name; write-view view-name; }</pre>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c)]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Define the security level used for access privileges.
Default	none
Options	<p>authentication—Provide authentication but no encryption.</p> <p>none—No authentication and no encryption.</p> <p>privacy—Provide authentication and encryption.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Security Level</i>

security-level (Generating SNMP Notifications)

Syntax	security-level (authentication none privacy);
Hierarchy Level	[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security level to use when generating SNMP notifications.
Default	none
Options	authentication —Provide authentication but no encryption. none —No authentication and no encryption. privacy —Provide authentication and encryption.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Level</i>

security-model (Access Privileges)

Syntax	<code>security-model (usm v1 v2c);</code>
Hierarchy Level	<code>[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>)]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security model for an SNMPv3 group. The security model is used to determine access privileges for the group.
Options	<code>usm</code> —SNMPv3 security model. <code>v1</code> —SNMPv1 security model. <code>v2c</code> —SNMPv2c security model.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Security Model</i>

security-model (Group)

Syntax	<pre>security-model (usm v1 v2c) { security-name security-name { group group-name; } }</pre>
Hierarchy Level	[edit snmp v3 vacm security-to-group]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Define a security model for an SNMPv3 group and associate the security name of a user with the group. All users in the group have the same access privileges.
Options	usm —SNMPv3 security model. v1 —SNMPv1 security model. v2c —SNMPv2c security model.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Model</i>


security-model (SNMP Notifications)

Syntax	<code>security-model (usm v1 v2c);</code>
Hierarchy Level	<code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security model for an SNMPv3 group. The security model is used for SNMP notifications.
Options	usm —SNMPv3 security model. v1 —SNMPv1 security model. v2c —SNMPv2c security model.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Security Model</i>


security-name (Security Group)

Syntax	<code>security-name security-name { group group-name; }</code>
Hierarchy Level	[edit snmp v3 vacm security-to-group security-model (usm v1 v2c)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate the security name of a user (for SNMPv3 clients) or a community string (for SNMPv1 and SNMPv2c clients) with a configured security group.
Options	security-name —SNMPv3 secure username configured at the [edit snmp v3 usm local-engine user username] hierarchy level that is used for messaging security. For SNMPv1 and SNMPv2c, the security name is the community string configured at the [edit snmp v3 snmp-community community-index] hierarchy level.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Assigning Security Names to Groups</i>• <i>Assigning a Security Name to a Group</i>

security-name (Community String)

Syntax	<code>security-name <i>security-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 <i>snmp-community</i> <i>community-index</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Associate a community string with the security name of a user. The community string, which is used for SNMPv1 and SNMPv2c clients in an SNMPv3 system, is configured at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.
Options	<i>security-name</i> —Name that is used for messaging security and user access control.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: The security name must match the configured security name at the <code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code> hierarchy level when you configure traps or informs.</p> </div> </div>	
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Security Names</i>

security-name (SNMP Notifications)

Syntax	<code>security-name <i>security-name</i>;</code>
Hierarchy Level	<code>[edit snmp v3 target-parameters <i>target-parameters-name</i> parameters]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the security name used when generating SNMP notifications.
Options	<i>security-name</i> —If the SNMPv3 USM security model is used, identify the user when generating the SNMP notification. If the v1 or v2c security models are used, identify the SNMP community used when generating the notification.
<div><div></div><div><p>NOTE: The access privileges for the group associated with this security name must allow this notification to be sent.</p><p>If you are using the v1 or v2 security models, the security name at the <code>[edit snmp v3 vacm security-to-group]</code> hierarchy level must match the security name at the <code>[edit snmp v3 snmp-community <i>community-index</i>]</code> hierarchy level.</p></div></div>	
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Security Name</i>

security-to-group

Syntax	<pre>security-to-group { security-model (usm v1 v2c) { group group-name; security-name security-name; } }</pre>
Hierarchy Level	[edit snmp v3 vacm]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Configure the group to which a specific SNMPv3 security name belongs. The security name is used for messaging security.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Assigning Security Model and Security Name to a Group</i>

snmp

Syntax	snmp { ... }
Hierarchy Level	[edit]
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>Configure SNMP.</p> <p>SNMP modules cannot have the slash (/) character or the @ character in the name.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring SNMP on a Device Running Junos OS</i>

snmp

Syntax

```
snmp {  
  rmon {  
    history index {  
      interface (SNMP RMON History) interface-name;  
      bucket-size number;  
      interval seconds;  
      owner owner-name;  
    }  
  }  
}
```

Hierarchy Level [\[edit\]](#)

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure SNMP.

The statements are explained separately.

Required Privilege Level

- snmp—To view this statement in the configuration.
- snmp-control—To add this statement to the configuration.

Related Documentation

- [Configuring SNMP \(J-Web Procedure\) on page 75](#)

snmp-community

Syntax	snmp-community <i>community-index</i> { community-name <i>community-name</i> ; security-name <i>security-name</i> ; tag <i>tag-name</i> ; }
Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the SNMP community which authorizes SNMPv1 or SNMPv2c clients in an SNMPv3 system.
Options	<i>community-index</i> —(Optional) String that identifies an SNMP community. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the SNMPv3 Community</i>

source-address (SNMP)

Syntax	source-address <i>address</i> ;
Hierarchy Level	[edit snmp trap-options]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Set the source address of every SNMP trap packet sent by this router to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.
Options	address —Source address of SNMP traps. You can configure the source address of trap packets two ways: lo0 or a valid IPv4 address configured on one of the router interfaces. The value lo0 indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface lo0 . Default: Disabled. (The source address is the address of the outgoing interface.)
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Source Address for SNMP Traps</i>

startup-alarm

Syntax	startup-alarm (falling-alarm rising-alarm rising-or-falling-alarm);
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	The alarm that can be sent upon entry startup.
Options	<p>falling-alarm—Generated if the first sample after the alarm entry becomes active is less than or equal to the falling threshold.</p> <p>rising-alarm—Generated if the first sample after the alarm entry becomes active is greater than or equal to the rising threshold.</p> <p>rising-or-falling-alarm—Generated if the first sample after the alarm entry becomes active satisfies either of the corresponding thresholds.</p> <p>Default: rising-or-falling-alarm</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Sample Type</i>

syslog-subtag

Syntax	syslog-subtag <i>syslog-subtag</i> ;
Hierarchy Level	[edit snmp rmon alarm index]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Add a tag to the system log message.
Options	<p>syslog-subtag <i>syslog-subtag</i>—Tag of not more than 80 uppercase characters to be added to syslog messages.</p> <p>Default: None</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the System Log Tag</i>

tag (SNMPv3)

Syntax	<code>tag tag-name;</code>
Hierarchy Level	[edit snmp v3 notify name], [edit snmp v3 snmp-community community-index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure a set of targets to receive traps or informs (for IPv4 packets only).
Options	tag-name —Identifies the address of managers that are allowed to use a community string.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Tag</i>• <i>Configuring the SNMPv3 Trap Notification</i>

tag-list

Syntax	<code>tag-list tag-list;</code>
Hierarchy Level	[edit snmp v3 target-address <i>target-address-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure an SNMP tag list used to select target addresses.
Options	tag-list —Define sets of target addresses (tags). To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Trap Target Address</i>

target-address

Syntax	<pre>target-address <i>target-address-name</i> { address <i>address</i>; address-mask <i>address-mask</i>; logical-system (SNMP) <i>logical-system</i>; port (SNMP) <i>port-number</i>; retry-count (SNMPv3) <i>number</i>; routing-instance (SNMPv3) <i>instance</i>; tag-list <i>tag-list</i>; target-parameters <i>target-parameters-name</i>; timeout <i>seconds</i>; }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	Configure the address of an SNMP management application and the parameters to be used in sending notifications.
Options	<p><i>target-address-name</i>—String that identifies the target address.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring the Trap Target Address

target-parameters

Syntax At the `[edit snmp v3]` hierarchy level:

```
target-parameters target-parameters-name {  
  profile-name;  
  parameters {  
    message-processing-model (v1 | v2c | V3);  
    security-level (authentication | none | privacy);  
    security-model (usm | v1 | v2c);  
    security-name security-name;  
  }  
}
```

At the `[edit snmp v3 target-address target-address-name]` hierarchy level:

```
target-parameters target-parameters-name;
```

Hierarchy Level `[edit snmp v3]`
`[edit snmp v3 target-address target-address-name]`

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.1 for the QFX Series.

Description Configure the message processing and security parameters for sending notifications to a particular management target. The target parameters are configured at the `[edit snmp v3]` hierarchy level. The remaining statements at this level are explained separately.

Then apply the target parameters configured at the `[edit snmp v3 target-parameters target-parameters-name]` hierarchy level to the target address configuration at the `[edit snmp v3]` hierarchy level.

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

Related Documentation

- *Defining and Configuring the Trap Target Parameters*
- *Applying Target Parameters*

targets

Syntax	<pre>targets { address; }</pre>
Hierarchy Level	[edit snmp trap-group <i>group-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	Configure one or more systems to receive SNMP traps.
Options	address —IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Groups</i>

traceoptions (SNMP)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; memory-trace; no-remote-trace; no-default-memory-trace; }</pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>file <i>filename</i> option added in Junos OS Release 8.1.</p> <p>world-readable no-world-readable option added in Junos OS Release 8.1.</p> <p>match <i>regular-expression</i> option added in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>memory-trace and no-default-memory-trace options introduced in Junos OS Release 13.3.</p>
Description	<p>The output of the tracing operations is placed into log files in the /var/log directory. Each log file is named after the SNMP agent that generates it. Currently, the following logs are created in the /var/log directory when the traceoptions statement is used:</p> <ul style="list-style-type: none">• chassisd• craftd• ilmids• mib2d• rmopd• serviced• snmpd
Options	<p>file <i>filename</i>—By default, the name of the log file that records trace output is the name of the process being traced (for example, mib2d or snmpd). Use this option to specify another name.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, snmpd) reaches its maximum size, it is archived by being renamed to snmpd.0. The previous snmpd.1 is renamed to snmpd.2, and so on. The oldest archived file is deleted.</p> <p>Range: 2 through 1000 files</p> <p>Default: 10 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements:</p>

- **all**—Log all SNMP events.
- **general**—Log general events.
- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **subagent**—Log subagent restarts.
- **timer**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

memory-trace—Enable tracing in memory. If **file** and **memory-trace** are configured, tracing will occur in file and memory.



NOTE: By default, memory tracing is enabled with default trace flags even after **snmp traceoptions** config statement is disabled.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-default-memory-trace—(Optional) Disable default memory-tracing.

size *size*—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

Range: 10 KB through 1 GB

Default: 1000 KB

world-readable | no-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

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

Related Documentation

- *Tracing SNMP Activity on a Device Running Junos OS*

trap-group

Syntax	<pre>trap-group group-name { categories { category; } destination-port port-number; routing-instance instance; targets { address; } version (all v1 v2); }</pre>
Hierarchy Level	[edit snmp]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.
Options	<p>group-name—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Trap Groups</i>

trap-options

Syntax	<pre>trap-options { agent-address outgoing-interface; source-address address; }</pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router or switch to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.</p> <p>The remaining statements are explained separately.</p>
Default	Disabled
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring SNMP Trap Options</i>

type (SNMP RMON)

Syntax	<code>type type;</code>
Hierarchy Level	[edit snmp rmon event index]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Type of notification generated when a threshold is crossed.
Options	<p>type—Type of notification:</p> <ul style="list-style-type: none">• log—Add an entry to logTable.• log-and-trap—Send an SNMP trap and make a log entry.• none—No notifications are sent.• snmptrap—Send an SNMP trap. <p>Default: log-and-trap</p>
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an RMON Event Entry and Its Attributes</i>

type (SNMPv3)

Syntax	type (inform trap);
Hierarchy Level	[edit snmp v3 notify <i>name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. inform option added in Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure the type of SNMP notification.
Options	inform —Defines the type of notification as an inform. SNMP informs are confirmed notifications. trap —Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Informs</i>• <i>Configuring the SNMPv3 Trap Notification</i>

v3

```

Syntax v3 {
    notify name {
        tag tag-name;
        type trap;
    }
    notify-filter profile-name {
        oid object-identifier (include | exclude);
    }
    snmp-community community-index {
        community-name community-name;
        security-name security-name;
        tag tag-name;
    }
    target-address target-address-name {
        address address;
        address-mask address-mask;
        logical-system (SNMP) logical-system;
        port port-number;
        retry-count number;
        routing-instance instance;
        tag-list tag-list;
        target-parameters target-parameters-name;
        timeout seconds;
    }
    target-parameters target-parameters-name {
        notify-filter profile-name;
        parameters {
            message-processing-model (v1 | v2c | V3);
            security-level (authentication | none | privacy);
            security-model (usm | v1 | v2c);
            security-name security-name;
        }
    }
    usm {
        local-engine {
            user username {
                authentication-md5 {
                    authentication-password authentication-password;
                }
                authentication-sha {
                    authentication-password authentication-password;
                }
                authentication-none;
                privacy-aes128 {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
                privacy-des {
                    privacy-password privacy-password;
                }
            }
        }
    }
}

```



```

        privacy-none;
    }
}
remote-engine engine-id {
    user username {
        authentication-md5 {
            authentication-password authentication-password;
        }
        authentication-sha {
            authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
            privacy-password privacy-password;
        }
        privacy-des {
            privacy-password privacy-password;
        }
        privacy-3des {
            privacy-password privacy-password;
        }
        privacy-none {
            privacy-password privacy-password;
        }
    }
}
}
}
vacm {
    access {
        group group-name {
            (default-context-prefix | context-prefix context-prefix) {
                security-model (any | usm | v1 | v2c) {
                    security-level (authentication | none | privacy) {
                        notify-view view-name;
                        read-view view-name;
                        write-view view-name;
                    }
                }
            }
        }
    }
}
}
security-to-group {
    security-model (usm | v1 | v2c) {
        security-name security-name {
            group group-name;
        }
    }
}
}
}
}

```

Hierarchy Level [edit snmp]

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

Description	Configure SNMPv3. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Minimum SNMPv3 Configuration on a Device Running Junos OS</i>

vacm

Syntax	<pre>vacm { access { group group-name { (default-context-prefix context-prefix context-prefix){ security-model (any usm v1 v2c) { security-level (authentication none privacy) { notify-view view-name; read-view view-name; write-view view-name; } } } } security-to-group { security-model (usm v1 v2c); security-name security-name { group group-name; } } } }</pre>
Hierarchy Level	[edit snmp v3]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Configure view-based access control model (VACM) information, including access privileges such as security model and security level for a group of users. The remaining statements are explained separately.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Defining Access Privileges for an SNMP Group</i>

variable

Syntax	<code>variable <i>oid-variable</i>;</code>
Hierarchy Level	<code>[edit snmp rmon alarm <i>index</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Object identifier (OID) of MIB variable to be monitored.
Options	<i>oid-variable</i> —OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, 1.3.6.1.2.1.2.1.2.1.10.1). Alternatively, use the MIB object name (for example, <code>ifInOctets.1</code>).
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring the Variable</i>


version (SNMP)

Syntax	<code>version (all v1 v2);</code>
Hierarchy Level	<code>[edit snmp trap-group <i>group-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Specify the version number of SNMP traps.
Default	<code>all</code> —Send an SNMPv1 and SNMPv2 trap for every trap condition.
Options	<code>all</code> —Send an SNMPv1 and SNMPv2 trap for every trap condition. <code>v1</code> —Send SNMPv1 traps only. <code>v2</code> —Send SNMPv2 traps only.
Required Privilege Level	<code>snmp</code> —To view this statement in the configuration. <code>snmp-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring SNMP Trap Groups</i>

view (Associating a MIB View with a Community)

Syntax	<code>view view-name;</code>
Hierarchy Level	<code>[edit snmp community community-name]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Associate a view with a community. A view represents a group of MIB objects.
Options	view-name —Name of the view. You must use a view name already configured in the view statement at the [edit snmp] hierarchy level.
Required Privilege Level	snmp —To view this statement in the configuration. snmp-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring SNMP Communities</i>

view (Configuring a MIB View)

Syntax	<pre>view view-name { oid object-identifier (include exclude); }</pre>
Hierarchy Level	[edit snmp]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.0 for EX Series switches.</p>
Description	<p>Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The view statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the view statement at the [edit snmp community community-name] hierarchy level.</p>
	<div>  <p>NOTE: To remove an OID completely, use the delete view all oid oid-number command but omit the include parameter.</p> </div>
Options	<p>view-name—Name of the view.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>snmp—To view this statement in the configuration.</p> <p>snmp-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring MIB Views</i> • <i>Associating MIB Views with an SNMP User Group</i> • community on page 121

write-view

Syntax	<code>write-view view-name;</code>
Hierarchy Level	[edit snmp v3 vacm access group <i>group-name</i> (default-context-prefix context-prefix <i>context-prefix</i>) security-model (any usm v1 v2c) security-level (authentication none privacy)]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series switches.
Description	Associate the write view with a community (for SNMPv1 or SNMPv2c clients) or a group name (for SNMPv3 clients).
Options	<i>view-name</i> —Name of the view for which the SNMP user group has write permission.
Required Privilege Level	snmp—To view this statement in the configuration. snmp-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring MIB Views</i>• <i>Configuring the Write View</i>

CHAPTER 9

Configuration Statements: Port Mirroring

- [\[edit ethernet-switching-options\] Configuration Statement Hierarchy on EX Series Switches](#) on page 185
- [analyzer](#) on page 189
- [egress](#) on page 190
- [egress \(Interface or VLAN\)](#) on page 190
- [ethernet-switching-options](#) on page 191
- [ingress \(vlans\)](#) on page 194
- [ingress \(Interface or VLAN\)](#) on page 195
- [input \(Port Mirroring\)](#) on page 196
- [interface](#) on page 197
- [loss-priority \(Port Mirroring\)](#) on page 198
- [no-tag](#) on page 199
- [output \(Port Mirroring\)](#) on page 200
- [ratio](#) on page 201
- [vlan \(Port Mirroring\)](#) on page 202

[\[edit ethernet-switching-options\] Configuration Statement Hierarchy on EX Series Switches](#)

This topic lists supported and unsupported configuration statements in the **[edit ethernet-switching-options]** hierarchy level on EX Series switches.

- *Supported* statements are those that you can use to configure some aspect of a software feature on the switch.
- *Unsupported* statements are those that appear in the command-line interface (CLI) on the switch, but that have no effect on switch operation if you configure them.
- Not all features are supported on all switch platforms. For detailed information about feature support on specific EX Series switch platforms, see [Feature Explorer](#).

This topic lists:

- [Supported Statements in the \[edit ethernet-switching-options\] Hierarchy Level on page 186](#)
- [Unsupported Statements in the \[edit ethernet-switching-options\] Hierarchy Level on page 188](#)

Supported Statements in the [edit ethernet-switching-options] Hierarchy Level

The following hierarchy shows the **[edit ethernet-switching-options]** configuration statements supported on EX Series switches:

```
ethernet-switching-options {
  analyzer {
    name {
      input {
        egress {
          interface (all | interface-name);
        }
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
      }
    }
    loss-priority priority;
    output {
      interface interface-name;
      vlan (vlan-id | vlan-name);
    }
    ratio number;
  }
}
authentication-whitelist {
  interface;
  vlan-assignment;
}
bpdu-block {
  disable-timeout timeout;
  interface (all | [interface-name]) {
    (disable | drop | shutdown);
  }
}
dot1q-tunneling {
  ether-type (0x8100 | 0x88a8 | 0x9100);
}
interfaces interface-name {
  no-mac-learning;
}
mac-lookup-length number-of-entries;
}
mac-notification {
  notification-interval seconds;
}
}
mac-table-aging-time seconds;
port-error-disable {
```



```

        disable-timeout timeout;
    }
    redundant-trunk-group {
        group name {
            description;
            interface interface-name {
                primary;
            }
            preempt-cutover-timer seconds;
        }
    }
    secure-access-port {
        dhcp-snooping-file {
            location local_pathname | remote_URL;
            timeout seconds;
            write-interval seconds;
        }
        interface (all | interface-name) {
            allowed-mac {
                mac-address-list;
            }
            (dhcp-trusted | no-dhcp-trusted );
            fcoe-trusted;
            mac-limit limit action action;
            no-allowed-mac-log;
            static-ip ip-address {
                mac mac-address;
                vlan vlan-name;
            }
        }
    }
    uac-policy;
}
vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection );
    dhcp-option82 {
        disable;
        circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
        }
        remote-id {
            prefix (hostname | mac | none);
            use-interface-description;
            use-string string;
        }
        vendor-id [string];
    }
    (examine-dhcp | no-examine-dhcp);
    examine-fip {
        fc-map fc-map-value;
    }
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
}
}

```

```
static {
  vlan vlan-id {
    mac mac-address next-hop interface-name;
  }
}
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    multicast;
    no-broadcast;
    no-multicast;
    no-registered-multicast;
    no-unknown-unicast;
    no-unregistered-multicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
}
unknown-unicast-forwarding {
  vlan (all | vlan-name) {
    interface interface-name;
  }
}
voip {
  interface (all | [interface-name | access-ports]) {
    forwarding-class (assured-forwarding | best-effort | expedited-forwarding |
      network-control);
    vlan vlan-name;
  }
}
}
```

Unsupported Statements in the [edit ethernet-switching-options] Hierarchy Level

All statements in the [edit ethernet-switching-options] hierarchy level that are displayed in the command-line interface (CLI) on the switch are supported on the switch and operate as documented.

Related Documentation

- *Example: Setting Up Q-in-Q Tunneling on EX Series Switches*
- *Example: Configuring Redundant Trunk Links for Faster Recovery*
- *Configuring MAC Table Aging (CLI Procedure)*
- *Configuring MAC Notification (CLI Procedure)*
- *Configuring Q-in-Q Tunneling (CLI Procedure)*
- *Configuring Redundant Trunk Links for Faster Recovery (CLI Procedure)*
- *Configuring Nonstop Bridging on EX Series Switches (CLI Procedure)*

analyzer

Syntax	<pre> analyzer { name { ratio number; loss-priority priority; input { ingress { interface (all interface-name); vlan (vlan-id vlan-name); } egress { interface (all interface-name); } } output { interface interface-name; vlan (vlan-id vlan-name) { no-tag; } } } } </pre>
Hierarchy Level	[edit ethernet-switching-options]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure port mirroring. One analyzer (port mirroring configuration) can be configured on an EX2200, EX3200, EX4200, or EX4500 switch and seven analyzers (port mirroring configurations) can be configured on an EX8208 or EX8216 switch at a time. Other analyzers can be present and disabled.
Default	Port mirroring is disabled and Junos OS creates no default analyzers.
Options	<p>name—Name that identifies the analyzer. The name can be up to 125 characters long, must begin with a letter, and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring Mirroring for Local Monitoring of Employee Resource Use on EX4300 Switches</i> • <i>Example: Configuring Mirroring for Remote Monitoring of Employee Resource Use on EX4300 Switches</i> • <i>Understanding Port Mirroring and Analyzers on EX4300 Switches</i>

egress

Syntax	<code>egress;</code>
Hierarchy Level	<code>[edit vlans <i>vlan-name</i> vlan-id <i>number</i> interface <i>interface-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Specify that the member interface of the VLAN allows only egress traffic.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61

egress (Interface or VLAN)

Syntax	<pre>egress { interface (all <i>interface-name</i>); vlan (<i>vlan-id</i> <i>vlan-name</i>); }</pre>
Hierarchy Level	<code>[edit ethernet-switching-options analyzer <i>name</i> input]</code>
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Specify ports or VLANs for which traffic exiting the interface or VLAN is mirrored in a port mirroring configuration. You can define the egress VLAN ID or VLAN name for port mirroring only for EX8200 switches.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61

ethernet-switching-options

```

Syntax  ethernet-switching-options {
        analyzer {
            name {
                loss-priority priority;
                ratio number;
            }
            input {
                ingress {
                    interface (all | interface-name);
                    vlan (vlan-id | vlan-name);
                }
                egress {
                    interface (all | interface-name);
                }
            }
            output {
                interface interface-name;
                vlan (vlan-id | vlan-name) {
                    no-tag;
                }
            }
        }
    }
    bpd-block {
        disable-timeout timeout;
        interface (all | [interface-name]) {
            (disable | drop | shutdown);
        }
    }
    dot1q-tunneling {
        ether-type (0x8100 | 0x88a8 | 0x9100);
    }
    interfaces interface-name {
        no-mac-learning;
    }
    mac-lookup-length number-of-entries;
}
mac-notification {
    notification-interval seconds;
}
mac-table-aging-time seconds;
nonstop-bridging;
port-error-disable {
    disable-timeout timeout;
}
redundant-trunk-group {
    group name {
        interface interface-name <primary>;
        interface interface-name;
    }
}
secure-access-port {
    dhcp-snooping-file {

```

```
location local_pathname | remote_URL;  
timeout seconds;  
write-interval seconds;  
}  
dhcpv6-snooping-file {  
    location local_pathname | remote_URL;  
    timeout seconds;  
    write-interval seconds;  
}  
interface (all | interface-name) {  
    allowed-mac {  
        mac-address-list;  
    }  
    (dhcp-trusted | no-dhcp-trusted);  
    fcoe-trusted;  
    mac-limit limit action (drop | log | none | shutdown);  
    no-allowed-mac-log;  
    persistent-learning;  
    static-ip ip-address {  
        vlan vlan-name;  
        mac mac-address;  
    }  
    static-ipv6 ip-address {  
        vlan vlan-name;  
        mac mac-address;  
    }  
}  
vlan (all | vlan-name) {  
    (arp-inspection | no-arp-inspection) [  
        forwarding-class class-name;  
    ]  
    dhcp-option82 {  
        circuit-id {  
            prefix hostname;  
            use-interface-description;  
            use-vlan-id;  
        }  
        remote-id {  
            prefix hostname | mac | none;  
            use-interface-description;  
            use-string string;  
        }  
        vendor-id [string];  
    }  
    (examine-dhcp | no-examine-dhcp) {  
        forwarding-class class-name;  
    }  
    (examine-dhcpv6 | no-examine-dhcpv6) {  
        forwarding-class class-name;  
    }  
    examine-fip {  
        fc-map fc-map-value;  
    }  
    (ip-source-guard | no-ip-source-guard);  
    (ipv6-source-guard | no-ipv6-source-guard);  
    mac-move-limit limit action (drop | log | none | shutdown);
```

```

    }
    (neighbor-discovery-inspection | no-neighbor-discovery-inspection);
no-option-37;
static {
    vlan name {
        mac mac-address {
            next-hop interface-name;
        }
    }
}
storm-control {
    action-shutdown;
    interface (all | interface-name) {
        bandwidth bandwidth;
        level level;
        multicast;
        no-broadcast;
        no-multicast;
        no-registered-multicast;
        no-unknown-unicast;
        no-unregistered-multicast;
    }
}
traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
    flag flag <disable>;
}
unknown-unicast-forwarding {
    vlan (all | vlan-name) {
        interface interface-name;
    }
}
}
voip {
    interface (all | [interface-name | access-ports]) {
        vlan vlan-name;
        forwarding-class (assured-forwarding | best-effort | expedited-forwarding |
            network-control);
    }
}
}
}

```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

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

Related Documentation	<ul style="list-style-type: none">• Understanding Port Mirroring on EX Series Switches on page 9• <i>Understanding Port Security</i>• <i>Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches</i>• <i>Understanding Redundant Trunk Links</i>• <i>Understanding Storm Control on EX Series Switches</i>• <i>Understanding 802.1X and VoIP on EX Series Switches</i>• <i>Understanding Q-in-Q Tunneling on EX Series Switches</i>• <i>Understanding Unknown Unicast Forwarding</i>• <i>Understanding MAC Notification on EX Series Switches</i>• <i>Understanding FIP Snooping</i>• <i>Understanding Nonstop Bridging on EX Series Switches</i>
------------------------------	---

ingress (vlans)

Syntax	ingress;
Hierarchy Level	[edit vlans <i>vlan-name</i> vlan-id <i>number</i> interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 10.0 for EX Series switches.
Description	Specify that the member interface of the VLAN allows only ingress traffic.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61

ingress (Interface or VLAN)

Syntax	<pre>ingress { interface (all <i>interface-name</i>); vlan (<i>vlan-id</i> <i>vlan-name</i>); }</pre>
Hierarchy Level	[edit ethernet-switching-options analyzer name input]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure ports or VLANs for which the entering traffic is mirrored as part of a port mirroring configuration. You can define the ingress VLAN ID or VLAN name for port mirroring for all switches except EX8200 switches.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61

input (Port Mirroring)

Syntax	<pre>input { ingress { interface (all <i>interface-name</i>); vlan (<i>vlan-id</i> <i>vlan-name</i>); } egress { interface (all <i>interface-name</i>); } }</pre>
Hierarchy Level	[edit ethernet-switching-options analyzer name]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Define the traffic to be mirrored in a port mirroring configuration—the definition can be a combination of:</p> <ul style="list-style-type: none">• Packets entering or exiting a port• Packets entering a VLAN on an EX2200, EX3200, EX4200, or EX4500 switch• Packets exiting a VLAN on an EX8200 switch <p>The remaining statements are explained separately.</p>
Default	No default.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42• Understanding Port Mirroring on EX Series Switches on page 9

interface

Syntax	interface (all <i>interface-name</i>);
Hierarchy Level	[edit ethernet-switching-options analyzer name input egress], [edit ethernet-switching-options analyzer name input ingress], [edit ethernet-switching-options analyzer name output]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure the interfaces for which traffic is mirrored.
Options	<p>all—Apply port mirroring to all interfaces on the switch. Mirroring a high volume of traffic can be performance intensive for the switch. Therefore, you should generally select specific input interfaces in preference to using the all keyword, or use the all keyword in combination with setting a ratio for statistical sampling.</p> <p><i>interface-name</i>—Apply port mirroring to the specified interface only.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37 • Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42 • Understanding Port Mirroring on EX Series Switches on page 9

loss-priority (Port Mirroring)

Syntax	<code>loss-priority <i>priority</i>;</code>
Hierarchy Level	[edit ethernet-switching-options analyzer name]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure a loss priority for mirrored packets. By default, the switch applies a lower priority to mirrored data than to regular port-to-port data—mirrored traffic is dropped in preference for regular traffic when capacity is exceeded. For port mirroring configurations with output to an analyzer VLAN, set the loss priority to high.
Default	Low
Options	<i>priority</i> —The value for priority can be low or high. Default: low
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Understanding Port Mirroring on EX Series Switches on page 9• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42

no-tag

Syntax	no-tag;
Hierarchy Level	<p>[edit ethernet-switching-options analyzer name output vlan (vlan-id vlan-name)]</p> <p>[edit forwarding-options port-mirroring family ethernet-switching output vlan (vlan-name vlan-id)]</p> <p>[edit forwarding-options port-mirroring instance instance-name family ethernet-switching output vlan (vlan-name vlan-id)]</p>
Release Information	<p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p>Hierarchy [edit forwarding-options port-mirroring family ethernet-switching output vlan] introduced in Junos OS Release 13.2.</p> <p>Hierarchy [edit forwarding-options port-mirroring instance instance-name family ethernet-switching output vlan] introduced in Junos OS Release 13.2.</p>
Description	Specify that remote port-mirroring packets are not tagged.
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Port Mirroring to Multiple Interfaces for Remote Monitoring of Employee Resource Use on EX Series Switches on page 51 • Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37 • Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42

output (Port Mirroring)

Syntax	<pre>output { interface <i>interface-name</i>; vlan (<i>vlan-id</i> <i>vlan-name</i>) { no-tag; } }</pre>
Hierarchy Level	[edit ethernet-switching-options analyzer name]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure the destination for mirrored traffic, either an interface on the switch, for local monitoring, or a VLAN, for remote monitoring. You can optionally configure the no-tag statement so that remote port mirroring packets are not tagged.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42

ratio

Syntax	<code>ratio number;</code>
Hierarchy Level	[edit ethernet-switching-options analyzer name]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Configure port mirroring to copy a sampling of packets, by setting a ratio of 1:x. A value of 1 for x mirrors every packet, and 2047 mirrors 1 out of every 2047 packets.</p> <p>On EX8200 switches, you can set a ratio only for ingress packets.</p> <p>On EX4500 switches, if you configure a ratio for any one of the port-based analyzers, that ratio automatically applies to all port-based analyzers.</p>
Default	1
Options	<p><i>number</i>—The number of packets in the sample, out of which 1 packet is mirrored.</p> <p>Range: 1 through 2047</p> <p>Default: 1</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>system-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Understanding Port Mirroring on EX Series Switches on page 9

vlan (Port Mirroring)

Syntax	<code>vlan (<i>vlan-id</i> <i>vlan-name</i>) { no-tag; }</code>
Hierarchy Level	[edit ethernet-switching-options analyzer <i>name</i> input egress], [edit ethernet-switching-options analyzer <i>name</i> input ingress], [edit ethernet-switching-options analyzer <i>name</i> output]
Release Information	Statement introduced in Junos OS Release 9.0 for EX Series switches.
Description	Configure mirrored traffic to be sent to a VLAN for remote monitoring. On a destination (output) VLAN, you can also configure the no-tag statement.
Options	<p><i>vlan-id</i>—Numeric VLAN identifier.</p> <p><i>vlan-name</i>—Name of the VLAN.</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use Through a Transit Switch on EX Series Switches on page 61• Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42

CHAPTER 10

Configuration Statements: sFlow Technology

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 203](#)
- [agent-id on page 205](#)
- [collector on page 205](#)
- [disable \(sFlow Monitoring Technology\) on page 206](#)
- [interfaces \(sFlow Monitoring Technology\) on page 207](#)
- [polling-interval on page 208](#)
- [sample-rate on page 209](#)
- [sflow on page 210](#)
- [source-ip on page 211](#)
- [udp-port on page 211](#)

[\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches](#)

Each of the following topics lists the statements at a subhierarchy of the **[edit protocols]** hierarchy:

- [\[edit protocols bfd\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols bgp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols connections\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols dcbx\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols dot1x\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols igmp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols igmp-snooping\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols isis\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols lacp\] Configuration Statement Hierarchy on EX Series Switches](#)
- [\[edit protocols link-management\] Configuration Statement Hierarchy on EX Series Switches](#)

- [*\[edit protocols lldp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols lldp-med\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols mld\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols mld-snooping\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols mpls\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols msdp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols mstp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols mvrp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols neighbor-discovery\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols oam\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols ospf\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols ospf3\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols pim\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols rip\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols ripng\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols router-advertisement\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols router-discovery\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols rstp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols rsvp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols sflow\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols stp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols uplink-failure-detection\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols vrrp\] Configuration Statement Hierarchy on EX Series Switches*](#)
- [*\[edit protocols vstp\] Configuration Statement Hierarchy on EX Series Switches*](#)

**Related
Documentation**

- [*EX Series Switch Software Features Overview*](#)
- [*EX Series Virtual Chassis Software Features Overview*](#)

agent-id

Syntax	<code>agent-id <i>ip-address</i>;</code>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the IP address to be assigned as the agent ID for the sFlow agent. By assigning an IP address, you ensure that the IP address is not dynamic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

collector

Syntax	<pre>collector { <i>ip-address</i>; udp-port <i>port-number</i>; }</pre>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	<p>Configure a remote collector for sFlow network traffic monitoring. The switch sends sFlow UDP datagrams to this collector for analysis. You can configure up to four collectors on the switch. You configure a collector by specifying its IP address and a UDP port.</p> <p>The remaining statements are explained separately.</p>
Default	The default port is 6343.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 203 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

disable (sFlow Monitoring Technology)

Syntax	disable;
Hierarchy Level	[edit protocols sflow], [edit protocols sflow interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	Disable the sFlow monitoring protocol on all interfaces on the switch or on the specified interface.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 203• Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33• Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

interfaces (sFlow Monitoring Technology)

Syntax	<pre> interfaces <i>interface-name</i> { polling-interval <i>seconds</i>; sample-rate { egress <i>number</i>; ingress <i>number</i>; } } </pre>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	<p>Configure sFlow network traffic monitoring on the specified interface on the switch. You can configure sFlow parameters such as polling interval and sampling rate with different values on different interfaces, and you can also disable sFlow monitoring on individual interfaces.</p> <p>The remaining statements are explained separately.</p>
Options	<i>interface-name</i> —Name of the interface on which to configure sFlow parameters.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 203 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

polling-interval

Syntax	<code>polling-interval seconds;</code>
Hierarchy Level	[edit protocols sflow], [edit protocols sflow interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	Configure the interval (in seconds) that the switch waits between port statistics update messages. <i>Polling</i> refers to the switch gathering various statistics for the network interfaces configured for sFlow monitoring and exporting the statistics to the configured sFlow collector.
Default	If no polling interval is configured for a particular interface, the switch waits the number of seconds that is configured as the polling interval in the global sFlow configuration. If no polling interval is specified in the global configuration, the switch waits 20 seconds between messages.
Options	<i>seconds</i> —Number of seconds between port statistics update messages. A value of 0 (zero) specifies that polling is disabled. Range: 0 through 3600 seconds Default: 20 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 203• Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84• Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33

sample-rate

Syntax	sample-rate { egress <i>number</i> ; ingress <i>number</i> ; }
Hierarchy Level	[edit protocols sflow], [edit protocols sflow interfaces <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches. The option <i>number</i> (which immediately followed sample-rate) is deprecated and options egress <i>number</i> and ingress <i>number</i> added in Junos OS Release 10.4 for EX Series switches.
Description	Specify the number of egress or ingress packets out of which one packet is sampled. If no interface sampling rates are configured, the global sampling rates take effect. If neither is configured, by default both ingress and egress packet sampling is disabled
Default	By default, both egress and ingress sampling rates are disabled.



NOTE: The **sample-rate *number*** (the global sampling rate) statement has been deprecated and might be removed from future product releases. We strongly recommend that you phase out its use.

Options	egress <i>number</i> —Value for egress sampling rate. Range: 100 through 1,073,741,823 ingress <i>number</i> —Value for ingress sampling rate. Range: 100 through 1,073,741,823
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 203 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33

sflow

Syntax sflow {
 agent-id *ip-address*;
 collector {
 ip-address;
 udp-port *port-number*;
 }
 interfaces *interface-name* {
 disable;
 polling-interval *seconds*;
 sample-rate {
 egress *number*;
 ingress *number*;
 }
 }
 polling-interval *seconds*;
 sample-rate {
 egress *number*;
 ingress *number*;
 }
 source-ip *ip-address*;
 }

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.3 for EX Series switches.
Options **agent-id** and **source-ip** added in Junos OS Release 10.2 for EX Series switches.

Description Configure sFlow technology to continuously monitor traffic at wire speed on specified interfaces simultaneously. sFlow data can be used to provide network traffic visibility information.

The remaining statements are explained separately.

Default sFlow technology is disabled by default.

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

Related Documentation

- [\[edit protocols\] Configuration Statement Hierarchy on EX Series Switches on page 203](#)
- [Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33](#)
- [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 84](#)

source-ip

Syntax	<code>source-ip <i>ip-address</i>;</code>
Hierarchy Level	[edit protocols sflow]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the IP address to be used for the sFlow datagrams. By configuring an IP address, you ensure that the IP address is not dynamic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

udp-port

Syntax	<code>udp-port <i>port-number</i>;</code>
Hierarchy Level	[edit protocols sflow collector]
Release Information	Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	Configure the UDP port for a remote collector for sFlow network traffic monitoring. The switch sends sFlow UDP datagrams to the collector for analysis.
Options	<p><i>port-number</i>—UDP port number for this collector.</p> <p>Default: 6343</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • [edit protocols] Configuration Statement Hierarchy on EX Series Switches on page 203 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

CHAPTER 11

Configuration Statements: Ethernet OAM Connectivity Fault Management

- [action-profile \(Applying to OAM CFM, for EX Series Switch Only\) on page 215](#)
- [age \(EX Series Switch Only\) on page 216](#)
- [auto-discovery \(EX Series Switch Only\) on page 216](#)
- [calculation-weight on page 217](#)
- [connectivity-fault-management \(EX Series Switch Only\) on page 218](#)
- [continuity-check \(EX Series Switch Only\) on page 219](#)
- [cycle-time on page 220](#)
- [data-tlv-size on page 221](#)
- [delay on page 222](#)
- [delay-variation on page 223](#)
- [direction \(EX Series Switch Only\) on page 224](#)
- [hold-interval \(OAM CFM, for EX Series Switch Only\) on page 224](#)
- [interface \(OAM CFM, for EX Series Switch Only\) on page 225](#)
- [interval \(EX Series Switch Only\) on page 226](#)
- [iteration-count on page 227](#)
- [iteration-period on page 227](#)
- [level \(EX Series Switch Only\) on page 228](#)
- [linktrace \(EX Series Switch Only\) on page 228](#)
- [loss-threshold \(EX Series Switch Only\) on page 229](#)
- [maintenance-association \(EX Series Switch Only\) on page 230](#)
- [maintenance-domain \(EX Series Switch Only\) on page 231](#)
- [measurement-type \(OAM LFM\) on page 232](#)
- [mep \(EX Series Switch Only\) on page 233](#)
- [mip-half-function \(EX Series Switch Only\) on page 234](#)
- [name-format \(EX Series Switch Only\) on page 235](#)
- [path-database-size \(EX Series Switch Only\) on page 235](#)

- [performance-monitoring \(OAM LFM\) on page 236](#)
- [priority \(Protocols OAM\) on page 237](#)
- [priority \(OAM Connectivity-Fault Management\) on page 237](#)
- [remote-mep \(EX Series Switch Only\) on page 238](#)
- [short-name-format on page 239](#)
- [sla-iterator-profile on page 240](#)
- [sla-iterator-profiles \(OAM LFM\) on page 241](#)
- [traceoptions \(OAM CFM, for EX Series Switch Only\) on page 242](#)

action-profile (Applying to OAM CFM, for EX Series Switch Only)

Syntax	<pre> action-profile <i>profile-name</i> { action { interface-down; } default-actions { interface-down; } event { adjacency-loss; } } </pre>
Hierarchy Level	<p>[edit protocols oam ethernet connectivity-fault-management]</p> <p>[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id remote-mep mep-id]</p>
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure a name and default action for an action profile.
Options	<p><i>profile-name</i>—Name of the action profile.</p> <p><i>action</i>—Defines the action to be taken when connectivity to the remote MEP is lost.</p> <p><i>interface-down</i>—Brings the interface down when a remote MEP connectivity failure is detected.</p> <p><i>default-actions</i>—Defines the default action to be taken when connectivity to the remote MEP is lost.</p> <p><i>interface-down</i>—Brings the interface down when a remote MEP connectivity failure is detected.</p> <p><i>event</i>—Defines the event to be monitored when a remote MEP connectivity failure is detected.</p> <p><i>adjacency-loss</i>—Defines the connectivity loss to the remote MEP.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93 • Junos OS Network Interfaces Configuration Guide

age (EX Series Switch Only)

Syntax	age (30m 10m 1m 30s 10s);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management linktrace]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.
Default	10 minutes
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

auto-discovery (EX Series Switch Only)

Syntax	auto-discovery;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Enable the MEP to accept continuity check messages from all remote MEPs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

calculation-weight

Syntax	calculation-weight { <code>delay</code> <i>delay-value</i> ; <code>delay-variation</code> <i>delay-variation-value</i> ; }
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure the calculation weight for delay and delay variation.



NOTE: This option is applicable only for two-way delay measurement.

The remaining statements are explained separately.

Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring an Iterator Profile • Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102 • delay on page 222 • delay-variation on page 223

connectivity-fault-management (EX Series Switch Only)

```
Syntax connectivity-fault-management {
    action-profile profile-name {
        action {
            interface-down;
        }
        default-actions {
            interface-down;
        }
        event {
            adjacency-loss;
        }
    }
    linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
    }
    maintenance-domain domain-name {
        level number;
        mip-half-function (none | default | explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
            continuity-check {
                hold-interval minutes;
                interface-status-tlv;
                interval (10m | 10s | 1m | 1s | 100ms);
                loss-threshold number;
                port-status-tlv;
            }
            mep mep-id {
                auto-discovery;
                direction down;
                interface interface-name;
                remote-mep mep-id {
                    action-profile profile-name;
                }
            }
        }
    }
    performance-monitoring {
        sla-iterator-profiles {
            profile-name {
                calculation-weight {
                    delay delay-value;
                    delay-variation delay-variation-value;
                }
                cycle-time cycle-time-value;
                iteration-period iteration-period-value;
                measurement-type two-way-delay;
                passive;
            }
        }
    }
}
```


}

Hierarchy Level	[edit protocols oam ethernet]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches. performance-monitoring introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93 • Junos OS Network Interfaces Configuration Guide


continuity-check (EX Series Switch Only)

Syntax	<pre>continuity-check { hold-interval <i>minutes</i>; interface-status-tlv; interval (10m 10s 1m 1s 100ms); loss-threshold <i>number</i>; port-status-tlv; }</pre>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Specify continuity check protocol options. The remaining statements are explained separately.
Options	interface-status-tlv —Includes interface status TLV in CCM. port-status-tlv —Includes port status TLV in CCM. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93 • Junos OS Network Interfaces Configuration Guide

cycle-time

Syntax	<code>cycle-time <i>cycle-time-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure the time (in milliseconds) taken between back-to-back transmissions of SLA frames for a single connection.
Options	<i>cycle-time-value</i> —Cycle time value in milliseconds. Range: 10 through 3,600,000 Default: 1000
Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring an Iterator Profile• Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102

data-tlv-size

Syntax	<code>data-tlv-size size;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i> remote-mep <i>remote-mep-id</i> sla-iterator-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	Configure the size of the data TLV portion of the Y.1731 data frame.
Options	<i>size</i> —Size of the data TLV portion of the Y.1731 data frame.
<div>  NOTE: This option is applicable only for two-way delay measurement. </div>	
Range: 1 through 1400 bytes	
Default: 1	
Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
Related Documentation	<ul style="list-style-type: none"> sla-iterator-profile on page 240 <i>Configuring a Remote MEP with an Iterator Profile</i>

delay

Syntax	<code>delay <i>delay-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> calculation-weight]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure the calculation weight for delay.
Options	<i>delay-value</i> —Calculation weight for delay.



NOTE: This option is applicable only for two-way delay measurement.

Range: 1 through 65,535

Default: 1

Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
---------------------------------	--

- | | |
|------------------------------|--|
| Related Documentation | <ul style="list-style-type: none">• Configuring an Iterator Profile• Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102• calculation-weight on page 217 |
|------------------------------|--|

delay-variation

Syntax	<code>delay-variation <i>delay-variation-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i> calculation-weight]
Release Information	Statement introduced in Junos OS Release 11.1. Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure the calculation weight for delay variation.
Options	<i>delay-variation-value</i> —Calculation weight for delay variation.



NOTE: This option is applicable only for two-way delay measurement.

Range: 1 through 65,535

Default: 1

Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
---------------------------------	--

Related Documentation	<ul style="list-style-type: none"> • Configuring an Iterator Profile • Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102 • calculation-weight on page 217
------------------------------	--

direction (EX Series Switch Only)

Syntax	direction down;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Specify that connectivity fault management (CFM) packets (CCMs) be transmitted only in one direction for the MEP, that is, the direction be set as down so that CCMs are transmitted only out of (not into) the interface configured on this MEP.
Options	down —Down MEP CCMs are transmitted only out (not into) of the interface configured on this MEP.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

hold-interval (OAM CFM, for EX Series Switch Only)

Syntax	hold-interval <i>minutes</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the time to wait before flushing the maintenance association end point (MEP) database, if no updates occur.
Options	minutes —Time to wait, in minutes.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

interface (OAM CFM, for EX Series Switch Only)

Syntax	<code>interface (interface-name ((ge- xe-) (fpc/pic/port fpc/pic/port.unit-number fpc/pic/port.unit-number vlan vlan-id)));</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure IEEE 802.1ag Operation, Administration, and Management (OAM) Connectivity Fault Management (CFM) support for the specified interface.
Options	interface-name —Interface to which the MEP is attached. It can be a physical Ethernet interface or a logical interface. If the interface is a trunk interface, the VLAN associated with the interface must have a VLAN ID.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

interval (EX Series Switch Only)

Syntax	interval (10m 10s 1m 1s 100ms 10ms);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> continuity-check]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the time between continuity check messages.
Options	<p>10m—10 minutes.</p> <p>10s—10 seconds.</p> <p>1m—1 minute.</p> <p>1s—1 second.</p> <p>100ms—100 milliseconds.</p> <p>10ms—10 milliseconds.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

iteration-count

Syntax	<code>iteration-count <i>count-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i> remote-mep <i>remote-mep-id</i> sla-iterator-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	Configure the number of iterations for which the connection partakes in the iterator for acquiring SLA measurements.
Options	<p><i>count-value</i>—Number of iterations for which the connection should partake in the iterator for acquiring SLA measurements.</p> <p>Range: 1 through 65,535</p> <p>Default: 0 (or infinite iterations)</p>
Required Privilege Level	<p>Configure—To enter configuration mode.</p> <p>Control—To modify any configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • sla-iterator-profile on page 240 • <i>Configuring a Remote MEP with an Iterator Profile</i>

iteration-period

Syntax	<code>iteration-period <i>iteration-period-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 11.1.</p> <p>Statement introduced in Junos OS Release 11.4 for EX Series switches.</p>
Description	Configure the iteration period, which is the maximum number of cycles per iteration (that is, the number of connections registered to an iterator cannot exceed this value).
Options	<p><i>iteration-period-value</i>—Maximum number of cycles per iteration.</p> <p>Range: 1 through 2000</p> <p>Default: 2000</p>
Required Privilege Level	<p>Configure—To enter configuration mode.</p> <p>Control—To modify any configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring an Iterator Profile</i> • Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102

level (EX Series Switch Only)

Syntax	<code>level <i>number</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i>]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure a number to be used in CFM messages to identify the maintenance association.
Options	<i>number</i> —Number used to identify the maintenance domain to which the CFM message belongs. Range: 0 through 7
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

linktrace (EX Series Switch Only)

Syntax	<pre>linktrace { age (30m 10m 1m 30s 10s); path-database-size <i>path-database-size</i>; }</pre>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure connectivity fault management linktrace parameters. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

loss-threshold (EX Series Switch Only)

Syntax	<code>loss-threshold <i>number</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the number of continuity check messages that can be lost before the remote MEP is marked as down.
Options	<i>number</i> —Number of continuity check messages that can be lost before the remote MEP is marked down.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

maintenance-association (EX Series Switch Only)

Syntax `maintenance-association ma-name {
 continuity-check {
 hold-interval minutes;
 interface-status-tlv;
 interval (10m | 10s | 1m | 1s| 100ms);
 loss-threshold number;
 port-status-tlv;
 }
 mep mep-id {
 auto-discovery;
 direction down;
 interface interface-name;
 remote-mep mep-id {
 action-profile profile-name;
 }
 }
 }
 }`

Hierarchy Level [edit protocols [oam ethernet connectivity-fault-management maintenance-domain](#)
 domain-name]

Release Information Statement introduced in Junos OS Release 10.2 for EX Series switches.

Description Configure the name of the maintenance association in IEEE-compliant format.

Options *ma-name*—The name of the maintenance association within the maintenance domain.

The remaining statements are explained separately.

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

Related • [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 93](#)
Documentation • [Junos OS Network Interfaces Configuration Guide](#)

maintenance-domain (EX Series Switch Only)

```
Syntax  maintenance-domain domain-name {
        level number;
        mip-half-function (none | default | explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
            continuity-check {
                hold-interval minutes;
                interface-status-tlv;
                interval (10m | 10s | 1m | 1s | 100ms);
                loss-threshold number;
                port-status-tlv;
            }
            mep mep-id {
                auto-discovery;
                direction down;
                interface interface-name;
                remote-mep mep-id {
                    action-profile profile-name;
                }
            }
        }
    }
```

Hierarchy Level [edit protocols [oam ethernet connectivity-fault-management](#)]

Release Information Statement introduced in Junos OS Release 10.2 for EX Series switches.

Description Configure the name of the maintenance domain in IEEE-compliant format.

Options *domain-name*—The name for the maintenance domain.

The remaining statements are explained separately.

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

Related Documentation

- [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 93](#)
- [Junos OS Network Interfaces Configuration Guide](#)

measurement-type (OAM LFM)

Syntax	measurement-type two-way-delay;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure the measurement type for the service-level agreement (SLA) frames. An SLA frame is a type of packet used to measure frame loss in Ethernet connections.
Options	two-way-delay —Use Y.1731-compliant two-way ETH-DM frames to measure frame loss.
Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements (CLI Procedure) on page 100• Configuring One-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure) on page 101• Configuring Two-Way Ethernet Frame Delay Measurements on Switches (CLI Procedure) on page 104

mep (EX Series Switch Only)

Syntax	<pre> mep <i>mep-id</i> { auto-discovery; direction down; interface <i>interface-name</i>; remote-mep <i>mep-id</i> { action-profile <i>profile-name</i>; } }</pre>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association ma-name]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Configure the numeric identifier of the maintenance association end point (MEP) within the maintenance association.
Options	<p>mep-id—Numeric identifier of the MEP.</p> <p>Range: 1 through 8191</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93 Junos OS Network Interfaces Configuration Guide

mip-half-function (EX Series Switch Only)

Syntax	mip-half-function (none default explicit);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i>]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Specify the OAM Ethernet CFM maintenance domain MIP half functions.



NOTE: Whenever a MIP is configured, the MIP half function value for all maintenance domains and maintenance associations must be the same.

Options	none —Specify to not use the mip-half-function. default —Specify to use the default mip-half-function. explicit —Specify an explicit mip-half-function.
Required Privilege Level	routing —To view this statement in the configuration. routing-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

name-format (EX Series Switch Only)

Syntax	name-format (character-string none dns mac+2oct);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Specify the format of the maintenance domain name.
Options	<p>character-string—The name is an ASCII character string.</p> <p>none—Name format none means that maintenance domain name is not used.</p> <p>dns—Name is in domain name service (DNS) format. For example: www.juniper.net.</p> <p>mac+2oct—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.</p> <p>Default: character-string</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93 • Junos OS Network Interfaces Configuration Guide

path-database-size (EX Series Switch Only)

Syntax	path-database-size <i>path-database-size</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management linktrace]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Specify the number of linktrace reply entries to be stored per linktrace request.
Options	<p>path-database-size—Database size (number of entries stored per request).</p> <p>Range: 1 through 500</p> <p>Default: 100</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93 • Junos OS Network Interfaces Configuration Guide

performance-monitoring (OAM LFM)

Syntax performance-monitoring {
 sla-iterator-profiles {
 profile-name {
 calculation-weight {
 delay delay-value;
 delay-variation delay-variation-value;
 }
 cycle-time cycle-time-value;
 iteration-period iteration-period-value;
 measurement-type two-way-delay;
 passive;
 }
 }
 }

Hierarchy Level [edit protocols [oam](#) [ethernet](#) [connectivity-fault-management](#)]

Release Information Statement introduced in Junos OS Release 11.4 for EX Series switches.

Description Specify performance monitoring support for Ethernet frame delay measurement.

The remaining statements are explained separately.

Required Privilege Level Configure—To enter configuration mode.
 Control—To modify any configuration.

Related Documentation

- [Configuring MEP Interfaces on Switches to Support Ethernet Frame Delay Measurements \(CLI Procedure\) on page 100](#)
- [Configuring One-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 101](#)
- [Configuring Two-Way Ethernet Frame Delay Measurements on Switches \(CLI Procedure\) on page 104](#)

priority (Protocols OAM)

Syntax	<code>priority <i>priority-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i> remote-mep <i>remote-mep-id</i> sla-iterator-profile <i>profile-name</i> sla-iterator-profile]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	Configure the priority of the iterator profile, which is the vlan-pcp value that is sent in the Y.1731 data frames.
Options	<p><i>priority-value</i>—Priority value, which is the vlan-pcp value that is sent in the Y.1731 data frames.</p> <p>Range: 0 through 7</p> <p>Default: 0</p>
Required Privilege Level	<p>Configure—To enter configuration mode.</p> <p>Control—To modify any configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • sla-iterator-profile on page 240 • <i>Configuring a Remote MEP with an Iterator Profile</i>

priority (OAM Connectivity-Fault Management)

Syntax	<code>priority <i>number</i>;</code>
Hierarchy Level	<p>[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i>]</p> <p>For EX Series Switches:</p> <p>[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i>]</p>
Release Information	Statement introduced in Junos OS Release 8.4.
Description	IEEE 802.1p priority bits used by the continuity check messages.
Options	<p><i>number</i>—Configure the IEEE 802.1p priority bits to be used in the VLAN header of the CFM packets.</p> <p>Range: 0 through 7</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"> • <i>Configuring a Maintenance Endpoint</i>

remote-mep (EX Series Switch Only)

Syntax	<code>remote-mep mep-id { action-profile profile-name; }</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]
Release Information	Statement introduced in Junos OS Release 10.2 for EX Series switches.
Description	Specify the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.
Options	<p>mep-id—Specify the numeric identifier of the MEP.</p> <p>Range: 1 through 8191</p> <p>The remaining statement is explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93• Junos OS Network Interfaces Configuration Guide

short-name-format

Syntax	short-name-format (character-string vlan 2octet rfc-2685-vpn-id);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name]
Release Information	Statement introduced in Junos OS Release 8.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport Routers.
Description	Specify the name format of the maintenance association name.
Options	<p>character-string—The name is an ASCII character string.</p> <p>vlan—The primary VLAN identifier.</p> <p>2octet—A number in the range 0 through 65,535.</p> <p>rfc-2685-vpn-id—A VPN identifier that complies with RFC 2685.</p> <p>Default: character-string</p>



NOTE: The PTX Series Packet Transport Routers support the vlan and 2octet options only.

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>Creating a Maintenance Association</i> • <i>Configuring Ethernet 802.1ag OAM on PTX Series Packet Transport Routers</i>

sla-iterator-profile

Syntax	<code>sla-iterator-profile <i>profile-name</i> { <i>data-tlv-size</i> <i>size</i>; <i>iteration-count</i> <i>count-value</i>; <i>priority</i> <i>priority-value</i>; }</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i> remote-mep <i>remote-mep-id</i>]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	Configure a remote MEP with an iterator profile and specify the options.
Options	<p><i>profile-name</i>—Name of the iterator profile configured for a remote MEP. For more information about configuring a remote MEP with an iterator profile, see <i>Configuring a Remote MEP with an Iterator Profile</i>.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	Configure—To enter configuration mode. Control—To modify any configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring an Iterator Profile</i>• <i>Configuring a Remote MEP with an Iterator Profile</i>• <i>Verifying the Configuration of an Iterator Profile</i>• <i>Managing Iterator Statistics</i>• <i>sla-iterator-profiles</i>

sla-iterator-profiles (OAM LFM)

Syntax	<pre>sla-iterator-profiles { profile-name { calculation-weight { delay delay-value; delay-variation delay-variation-value; } cycle-time cycle-time-value; iteration-period iteration-period-value; measurement-type two-way-delay; passive; } }</pre>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
Release Information	Statement introduced in Junos OS Release 11.4 for EX Series switches.
Description	Configure an iterator application and specify the iterator profile options.
Options	<p><i>profile-name</i>—Name of the iterator profile.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>Configure—To enter configuration mode.</p> <p>Control—To modify any configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102

traceoptions (OAM CFM, for EX Series Switch Only)

Syntax `traceoptions {
 file filename <files number> <match regex> <size size> <world-readable |
 no-world-readable>;
 flag flag ;
 no-remote-trace;
 }`

Release Information Statement introduced in JUNOS OS Release 10.2 for EX Series switches.

Description Configure tracing options the connectivity fault management.

Options **file *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**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—Trace all events.
- **configuration**—Trace configuration events.
- **error**—Trace events related to catastrophic errors in daemon.
- **init**—Trace events related to protocol daemon start-up.
- **protocol**—Trace protocol processing events.
- **routing socket**—Trace routing socket events.

match—(Optional) Refine the output to log only those lines that match the given regular expression.

no-world-readable—(Optional) Restrict file access to the user who created the file.

no-remote-trace—(Optional) Disable the remote trace.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

Default: If you do not include this option, tracing output is appended to an existing trace file.

world-readable—(Optional) Enable unrestricted file access.

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

Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93
------------------------------	---

CHAPTER 12

Configuration Statements: Ethernet OAM Link Fault Management

- [action \(OAM LFM\) on page 246](#)
- [action-profile on page 247](#)
- [allow-remote-loopback on page 248](#)
- [ethernet \(OAM LFM\) on page 249](#)
- [event \(OAM LFM\) on page 252](#)
- [event-thresholds on page 252](#)
- [frame-error on page 253](#)
- [frame-period on page 253](#)
- [frame-period-summary on page 254](#)
- [interface \(OAM LFM\) on page 255](#)
- [link-adjacency-loss on page 256](#)
- [link-discovery on page 256](#)
- [link-down on page 257](#)
- [link-event-rate on page 257](#)
- [link-fault-management on page 258](#)
- [negotiation-options on page 259](#)
- [no-allow-link-events on page 259](#)
- [oam on page 260](#)
- [pdu-interval on page 262](#)
- [pdu-threshold on page 263](#)
- [remote-loopback on page 263](#)
- [symbol-period on page 264](#)
- [syslog \(OAM LFM\) on page 264](#)
- [traceoptions \(OAM LFM\) on page 265](#)

action (OAM LFM)

Syntax	<pre>action { syslog; link-down; }</pre>
Hierarchy Level	[edit protocols oam ethernet link-fault-management]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	<p>Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

action-profile

Syntax	<pre> action-profile <i>profile-name</i>; action { syslog; link-down; } event { link-adjacency-loss; link-event-rate { frame-error <i>count</i>; frame-period <i>count</i>; frame-period-summary <i>count</i>; symbol-period <i>count</i>; } } </pre>
Hierarchy Level	[edit protocols oam ethernet link-fault-management]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	<p>Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.</p> <p>The remaining statements are explained separately.</p>
Options	<i>profile-name</i> —Name of the action profile.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

allow-remote-loopback

Syntax	allow-remote-loopback;
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the EX Series switch.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

ethernet (OAM LFM)

```
Syntax  ethernet {
    connectivity-fault-management {
        action-profile profile-name {
            action {
                interface-down;
            }
            default-actions {
                interface-down;
            }
            event {
                adjacency-loss;
            }
        }
    }
    esp-traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable |
            no-world-readable>;
        flag (all | error | esp | interface | krt | lib | normal | task | timer);
    }
    linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
    }
    maintenance-domain domain-name {
        level number;
        mip-half-function (none | default | explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
            continuity-check {
                hold-interval minutes;
                interface-status-tlv;
                interval (10m | 10s | 1m | 1s | 100ms);
                loss-threshold number;
                port-status-tlv;
            }
            mep mep-id {
                auto-discovery;
                direction down;
                interface interface-name;
                priority
                remote-mep mep-id {
                    action-profile profile-name;
                    sla-iterator-profile profile-name {
                        data-tlv-size size;
                        iteration-count count-value;
                        priority priority-value;
                    }
                }
            }
        }
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
    }
    performance-monitoring {
```

```

sla-iterator-profiles {
  profile-name {
    calculation-weight {
      delay delay-value;
      delay-variation delay-variation-value;
    }
    cycle-time cycle-time-value;
    iteration-period iteration-period-value;
    measurement-type two-way-delay;
    passive;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag ;
  no-remote-trace;
}

```



```
    }  
  }  
}
```

Hierarchy Level	[edit protocols oam]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches. connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.
Description	<p>Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72• Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 68• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97• Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 93

event (OAM LFM)

Syntax	<pre>event { link-adjacency-loss; link-event-rate { frame-error count; frame-period count; frame-period-summary count; symbol-period count; } }</pre>
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure link events in an action profile for Ethernet OAM link fault management (LFM). The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

event-thresholds

Syntax	<pre>event-thresholds { frame-error count; frame-period count; frame-period-summary count; symbol-period count; }</pre>
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure threshold limit values for link events in periodic OAM PDUs. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

frame-error

Syntax	<code>frame-error count;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management event link-event-rate], [edit protocols oam ethernet link-fault-management interface interface-name event-thresholds]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure the threshold value for sending frame error events or taking the action specified in the action profile. Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.
Options	<i>count</i> —Threshold count in seconds for frame error events. Range: 1 through 100 seconds Default: 1 second
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

frame-period

Syntax	<code>frame-period count;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management event link-event-rate], [edit protocols oam ethernet link-fault-management interface interface-name event-thresholds]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure the number of frame errors within the last N frames that has exceeded a threshold. Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.
Options	<i>count</i> —Threshold count in seconds for frame error events. Range: 1 through 100 seconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

frame-period-summary

Syntax	<code>frame-period-summary count;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management event link-event-rate], [edit protocols oam ethernet link-fault-management interface interface-name event-thresholds]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	<p>Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period.</p>
Options	<p><i>count</i>—Threshold count in seconds for frame period summary error events.</p> <p>Range: 1 through 100 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

interface (OAM LFM)

Syntax	<pre> interface <i>interface-name</i> { link-discovery (active passive); pdu-interval <i>interval</i>; pdu-threshold <i>threshold-value</i>; remote-loopback; event-thresholds { frame-error <i>count</i>; frame-period <i>count</i>; frame-period-summary <i>count</i>; symbol-period <i>count</i>; } negotiation-options { allow-remote-loopback; no-allow-link-events; } } </pre>
Hierarchy Level	[edit protocols oam ethernet link-fault-management]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	<p>Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.</p> <p>The remaining statements are explained separately.</p>
Options	<i>interface-name</i> —Name of the interface to be enabled for IEEE 802.3ah OAM link fault management (LFM) support.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72 • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

link-adjacency-loss

Syntax	link-adjacency-loss;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure loss of adjacency event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the action statement.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

link-discovery

Syntax	link-discovery (active passive);
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface interface-name]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs.
Options	<p><i>active</i>—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.</p> <p><i>passive</i>—In passive mode, the peer initiates the discovery process.</p> <p>Once the discovery process is initiated, both sides participate in discovery.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

link-down

Syntax	link-down;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile action]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Mark the interface as down for transit traffic.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

link-event-rate

Syntax	link-event-rate { frame-error <i>count</i> ; frame-period <i>count</i> ; frame-period-summary <i>count</i> ; symbol-period <i>count</i> ; }
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure the number of link fault management (LFM) events per second. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

link-fault-management

```
Syntax  link-fault-management {
        action-profile profile-name;
        action {
            syslog;
            link-down;
        }
        event {
            link-adjacency-loss;
            link-event-rate {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
        }
        interface interface-name {
            link-discovery (active | passive);
            pdu-interval interval;
            pdu-threshold threshold-value;
            remote-loopback;
            event-thresholds {
                frame-error count;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
            }
            negotiation-options {
                allow-remote-loopback;
                no-allow-link-events;
            }
        }
    }
```

Hierarchy Level [edit protocols [oam](#) [ethernet](#)]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.

Description Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

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

Related Documentation

- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72](#)
- [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 97](#)

negotiation-options

Syntax	negotiation-options { allow-remote-loopback; no-allow-link-events; }
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces. The remaining statements are explained separately.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

no-allow-link-events

Syntax	no-allow-link-events;
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i> negotiation-options]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Disable the sending of link event TLVs.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

oam

```

Syntax  oam {
    ethernet {
        connectivity-fault-management {
            action-profile profile-name {
                action {
                    interface-down;
                }
                default-actions {
                    interface-down;
                }
                event {
                    adjacency-loss;
                }
            }
        }
        linktrace {
            age (30m | 10m | 1m | 30s | 10s);
            path-database-size path-database-size;
        }
        maintenance-domain domain-name {
            level number;
            mip-half-function (none | default | explicit);
            name-format (character-string | none | dns | mac+2oct);
            maintenance-association ma-name {
                continuity-check {
                    hold-interval minutes;
                    interface-status-tlv;
                    interval (10m | 10s | 1m | 1s | 100ms);
                    loss-threshold number;
                    port-status-tlv;
                }
                mep mep-id {
                    auto-discovery;
                    direction down;
                    interface interface-name;
                    remote-mep mep-id {
                        action-profile profile-name;
                    }
                }
            }
        }
    }
    performance-monitoring {
        sla-iterator-profiles {
            profile-name {
                calculation-weight {
                    delay delay-value;
                    delay-variation delay-variation-value;
                }
                cycle-time cycle-time-value;
                iteration-period iteration-period-value;
                measurement-type two-way-delay;
                passive;
            }
        }
    }
}

```

```

    }
  }
}
link-fault-management {
  action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
  }
}
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  pdu-threshold threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
}
}
}

```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 9.4 for EX Series switches.
connectivity-fault-management introduced in Junos OS Release 10.2 for EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on EX Series switches or configure connectivity fault management (CFM) for IEEE 802.lag Operation, Administration, and Management (OAM) support on the switches.

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**
- [Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72](#)
 - [Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches on page 68](#)
 - [Configuring Ethernet OAM Link Fault Management \(CLI Procedure\) on page 97](#)
 - [Configuring Ethernet OAM Connectivity Fault Management \(CLI Procedure\) on page 93](#)

pdu-interval

Syntax	<code>pdu-interval <i>interval</i>;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Specify the periodic OAM PDU sending interval for fault detection. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.
Options	<i>interval</i> —Periodic OAM PDU sending interval. Range: 400 through 1000 milliseconds Default: 1000 milliseconds
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

pdu-threshold

Syntax	<code>pdu-threshold <i>threshold-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.
Options	<p><i>threshold-value</i> —Number of PDUs missed before declaring the peer lost.</p> <p>Range: 3 through 10 PDUs</p> <p>Default: 3 PDUs</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

remote-loopback

Syntax	<code>remote-loopback;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72 • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

symbol-period

Syntax	<code>symbol-period count;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile <i>profile-name</i> ; event link-event-rate] , [edit protocols oam ethernet link-fault-management interface <i>interface-name</i> event-thresholds]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	<p>Configure the threshold for sending symbol period events or taking the action specified in the action profile.</p> <p>Symbol code errors occur on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period. You cannot configure the default value to a different value.</p>
Options	<p><i>count</i>—Threshold count in seconds for symbol period events.</p> <p>Range: 1 through 100 seconds</p>
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

syslog (OAM LFM)

Syntax	<code>syslog;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile <i>profile-name</i> ; action]
Release Information	Statement introduced in Junos OS Release 9.4 for EX Series switches.
Description	Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97

traceoptions (OAM LFM)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regex</i>> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i> ; no-remote-trace; } </pre>
Release Information	Statement introduced in JUNOS Release 10.2 for EX Series switches.
Description	Configure tracing options the link fault management.
Options	<p>file <i>filename</i>—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.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum xk to specify KB, xm to specify MB, or xg to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i>—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • action-profile—Trace action profile invocation events. • all—Trace all events. • configuration—Trace configuration events. • protocol—Trace protocol processing events. • routing socket—Trace routing socket events. <p>match—(Optional) Refine the output to log only those lines that match the given regular expression.</p> <p>no-world-readable—(Optional) Restrict file access to the user who created the file.</p> <p>no-remote-trace—(Optional) Disable the remote trace.</p> <p>size <i>size</i>—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option.</p> <p>Syntax: xk to specify KB, xm to specify MB, or xg to specify GB</p>

Range: 10 KB through 1 GB

Default: 128 KB

Default: If you do not include this option, tracing output is appended to an existing trace file.

world-readable—(Optional) Enable unrestricted file access.

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

Related Documentation	<ul style="list-style-type: none">• Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72• Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97
------------------------------	--

CHAPTER 13

Configuration Statements: RPM

- [data-fill on page 268](#)
- [data-size on page 269](#)
- [destination-port on page 270](#)
- [dscp-code-point on page 271](#)
- [hardware-timestamp on page 272](#)
- [history-size on page 273](#)
- [moving-average-size on page 273](#)
- [one-way-hardware-timestamp on page 274](#)
- [port \(RPM\) on page 274](#)
- [probe on page 275](#)
- [probe-count on page 276](#)
- [probe-interval on page 276](#)
- [probe-limit on page 277](#)
- [probe-server on page 278](#)
- [probe-type on page 279](#)
- [routing-instance on page 280](#)
- [routing-instances on page 280](#)
- [rpm \(Interfaces\) on page 281](#)
- [source-address \(Services\) on page 281](#)
- [tcp on page 282](#)
- [test on page 283](#)
- [test-interval on page 284](#)
- [thresholds on page 285](#)
- [traps on page 286](#)
- [udp on page 287](#)

data-fill

Syntax	<code>data-fill data;</code>
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 9.3. for PTX Series Packet Transport Routers.
Description	Specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes. The data-fill statement is not valid with the http-get or http-metadata-get probe types.
Options	data —A hexadecimal value; for example, 0-9 , A-F .
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>

data-size

Syntax	<code>data-size size;</code>
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 size of the data portion of ICMP probes. The data-size statement is not valid with the http-get or http-metadata-get probe type.
Options	data —The size can be from 0 through 65400 Default: 0



NOTE: If you configure the hardware timestamp feature (see *Configuring RPM Timestamping*):

- The **data-size** default value is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.
- The **data-size** must be at least 100 bytes smaller than the default MTU of the interface of the RPM client 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"> • <i>Configuring BGP Neighbor Discovery Through RPM</i>

destination-port

Syntax	<code>destination-port <i>port</i>;</code>
Hierarchy Level	[edit services rpm bgp], [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 Series Packet Transport Routers.
Description	<p>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.</p> <p>The value for the destination-port can be only 7 when you configure along with hardware timestamping. A constraint check prevents you for configuring any other value for the destination port in this case.</p> <p>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.</p>
Options	port —The port number can be 7 or from 49,160 to 65,535.
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>

dscp-code-point

Syntax	<code>dscp-code-point <i>dscp-bits</i>;</code>
Hierarchy Level	[edit services rpm probe <i>owner</i> 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 for PTX Series Packet Transport Routers.
Description	Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.
Options	<p><i>dscp-bits</i>—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:</p> <ul style="list-style-type: none"> • af11—Default: 001010 • af12—Default: 001100 • af13—Default: 001110 • af21—Default: 010010 • af22—Default: 010100 • af23—Default: 010110 • af31—Default: 011010 • af32—Default: 011100 • af33—Default: 011110 • af41—Default: 100010 • af42—Default: 100100 • af43—Default: 100110 • be—Default: 000000 • cs1—Default: 001000 • cs2—Default: 010000 • cs3—Default: 011000 • cs4—Default: 100000 • cs5—Default: 101000 • cs6—Default: 110000 • cs7—Default: 111000 • ef—Default: 101110 • nc1—Default: 110000

- **nc2**—Default: 111000

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

Related Documentation

- *Configuring RPM Probes*

hardware-timestamp

Syntax hardware-timestamp;

Hierarchy Level [edit services rpm **probe** owner **test** test-name]

Release Information Statement introduced in Junos OS Release 8.1.
Statement applied to MX Series routers in Junos OS Release 10.0.
Statement introduced in Junos OS Release 10.3 for EX Series switches.

Description On MX Series routers, on M-320 routers using the Enhanced Queuing MPC, and on EX Series switches only, enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with **icmp-ping**, **icmp-ping-timestamp**, **udp-ping**, and **udp-ping-timestamp** probe types.

When you configure either **probe-type udp-ping** or **probe-type udp-ping-timestamp** along with the **hardware-timestamp** command, the value for the **destination-port** can be only 7. 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 configuring **one-way-hardware-timestamp**.

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

Related Documentation

- *Configuring RPM Timestamping*

history-size

Syntax	<code>history-size size;</code>
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 number of stored history entries.
Options	size —A value from 0 to 512. Default: 50
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 BGP Neighbor Discovery Through RPM</i> • <i>Configuring RPM Probes</i>

moving-average-size

Syntax	<code>moving-average-size number;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test test-name]
Release Information	Statement introduced in Junos OS Release 8.5. 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	Enable statistical calculation operations to be performed across a configurable number of the most recent samples.
Options	number —Number of samples to be used in calculations. Range: 0 through 255
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>

one-way-hardware-timestamp

Syntax	<code>one-way-hardware-timestamp;</code>
Hierarchy Level	[edit services rpm probe owner test test-name]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	Enable timestamping of RPM probe messages for one-way delay and jitter measurements. You must configure this statement along with the destination-interface statement to invoke timestamping. This feature is supported only with icmp-ping , icmp-ping-timestamp , udp-ping , and udp-ping-timestamp probe types.
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 Timestamping</i>• <i>destination-interface</i>• hardware-timestamp on page 272

port (RPM)

Syntax	<code>port number;</code>
Hierarchy Level	[edit services rpm probe-server (tcp udp)]
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 port number for the probe server.
Options	number —Port number for the probe server. The value can be 7 or 49,160 through 65,535.
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 Receiver Servers</i>

probe

Syntax	<pre> 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; moving-average-size number; one-way-hardware-timestamp; probe-count count; probe-interval seconds; probe-type type; routing-instance instance-name; source-address address; target (url address); test-interval interval; thresholds thresholds; traps traps; } } </pre>
Hierarchy Level	[edit services rpm]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p>
Description	Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.
Options	<p>owner—Specify an owner name up to 32 characters in length.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>system—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 RPM Probes

probe-count

Syntax	<code>probe-count count;</code>
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 number of probes within a test.
Options	count —A value from 1 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 BGP Neighbor Discovery Through RPM</i>• <i>Configuring RPM Probes</i>

probe-interval

Syntax	<code>probe-interval interval;</code>
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 time to wait between sending packets, in seconds.
Options	interval —Number of seconds, from 1 through 255.
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 BGP Neighbor Discovery Through RPM</i>• <i>Configuring RPM Probes</i>

probe-limit

Syntax	<code>probe-limit <i>limit</i>;</code>
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	limit —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

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

Syntax	<code>probe-type type;</code>
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>

routing-instance

Syntax	<code>routing-instance <i>instance-name</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code>
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 routing instance used by the probes.
Options	<i>instance-name</i> —A routing instance configured at the <code>[edit routing-instance]</code> hierarchy level. Default: Internet routing table <code>inet.0</code> .
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Probes</i>

routing-instances

Syntax	<code>routing-instances <i>instance-name</i>;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit services rpm bgp logical-system <i>logical-system-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.6. 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 routing instance used by the probes.
Options	<i>instance-name</i> —A routing instance configured at the <code>[edit routing-instances]</code> hierarchy level. Default: Internet routing table <code>inet.0</code> .
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery Through RPM</i>

rpm (Interfaces)

Syntax	<code>rpm (client server);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in Junos OS Release 8.1. Statement introduced in Junos OS Release 9.3 for EX Series switches.
Description	Associate an RPM client (router or switch that originates RPM probes) or RPM server with a specified interface.
Options	<i>client</i> —Identifier for RPM client router or switch. <i>server</i> —Identifier for RPM server.
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 Timestamping</i>

source-address (Services)

Syntax	<code>source-address <i>address</i>;</code>
Hierarchy Level	[edit services rpm <i>probe owner test 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 Series Packet Transport Routers.
Description	Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet will use the outgoing interface's address as its source.
Options	<i>address</i> —Valid IP address.
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>

tcp

Syntax	<pre>tcp { destination-interface <i>interface-name</i>; port <i>port</i>; }</pre>
Hierarchy Level	[edit services rpm probe-server]
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 port information for the TCP server. 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	<ul style="list-style-type: none">• <i>Configuring RPM Receiver Servers</i>

test

Syntax	<pre>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; moving-average-size number; one-way-hardware-timestamp; probe-count count; probe-interval seconds; probe-type type; routing-instance instance-name; source-address address; target (url url address address); test-interval interval; thresholds thresholds; traps traps; }</pre>
Hierarchy Level	[edit services rpm probe owner]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p>
Description	Specify the range of probes over which the standard deviation, average, and jitter are calculated. The test name combined with the owner name represent a single RPM configuration instance.
Options	<p>test-name—Specify a test name. The name can be up to 32 characters in length.</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 RPM Probes

test-interval

Syntax	<code>test-interval frequency;</code>
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 time to wait between tests, in seconds.
Options	frequency —Number of seconds, from 0 through 86400.
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 BGP Neighbor Discovery Through RPM</i>• <i>Configuring RPM Probes</i>

thresholds

Syntax	<code>thresholds thresholds;</code>
Hierarchy Level	[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 Packet Series Transport 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.
Options	<p>thresholds—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	<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 RPM Probes</i>

traps

Syntax	<code>traps traps;</code>
Hierarchy Level	[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	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">• egress-jitter-exceeded—Generates traps when the jitter in egress time threshold is met or exceeded.• egress-std-dev-exceeded—Generates traps when the egress time standard deviation threshold is met or exceeded.• egress-time-exceeded—Generates traps when the maximum egress time threshold is met or exceeded.• ingress-jitter-exceeded—Generates traps when the jitter in ingress time threshold is met or exceeded.• ingress-std-dev-exceeded—Generates traps when the ingress time standard deviation threshold is met or exceeded.• ingress-time-exceeded—Generates traps when the maximum ingress time threshold is met or exceeded.• jitter-exceeded—Generates traps when the jitter in round-trip time threshold is met or exceeded.• probe-failure—Generates traps for successive probe loss thresholds crossed.• rtt-exceeded—Generates traps when the maximum round-trip time threshold is met or exceeded.• std-dev-exceeded—Generates traps when the round-trip time standard deviation threshold is met or exceeded.• test-completion—Generates traps when a test is completed.• test-failure—Generates traps when the total probe loss threshold is met or exceeded.



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

- *Configuring RPM Probes*
- [categories on page 118](#)

udp

Syntax

```
udp {
    destination-interface interface-name;
    port port;
}
```

Hierarchy Level [edit services rpm [probe-server](#)]

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 port information for the UDP server.
 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*

CHAPTER 14

Configuration Statements: Uplink Failure Detection

- [action \(Uplink Failure Detection\)](#) on page 289
- [group \(Uplink Failure Detection\)](#) on page 290
- [link-to-disable](#) on page 290
- [link-to-monitor](#) on page 291
- [traceoptions \(Uplink Failure Detection\)](#) on page 292
- [uplink-failure-detection](#) on page 293

[action \(Uplink Failure Detection\)](#)

Syntax	<pre>action { log; }</pre>
Hierarchy Level	[edit protocols uplink-failure-detection]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Define an action on uplink-failure-detection group state change.
Options	log —Generate a system log message.
Required Privilege Level	admin —To view this statement in the configuration. admin-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 99

group (Uplink Failure Detection)

Syntax	<code>group <i>group-name</i> { link-to-monitor <i>interface-name</i>; link-to-disable <i>interface-name</i>; }</code>
Hierarchy Level	[edit protocols uplink-failure-detection]
Release Information	Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	Configure a group of uplink and downlink interfaces for uplink failure detection.
Options	<i>group-name</i> —Name of the uplink-failure-detection group. The remaining statements are explained separately.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 99

link-to-disable

Syntax	<code>link-to-disable <i>interface-name</i>;</code>
Hierarchy Level	[edit protocols uplink-failure-detection group <i>group-name</i>]
Release Information	Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	Configure the downlink interfaces to be disabled when the switch detects an uplink failure. The switch can monitor a maximum of 48 downlink interfaces in a group.
Options	<i>interface-name</i> —Name of the downlink interface or interface range in the group. The interface can be a physical interface or a logical interface.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 99

link-to-monitor

Syntax	<code>link-to-monitor <i>interface-name</i>;</code>
Hierarchy Level	<code>[edit protocols uplink-failure-detection group <i>group-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 11.1 for EX Series switches.
Description	<p>Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of 48 uplink interfaces in a group.</p> <p>An interface can be configured as link-to-monitor in multiple groups.</p>
Options	<i>interface-name</i> —Name of the uplink interface or interface range in the group. The interface can be a physical interface or a logical interface.
Required Privilege Level	<p>admin—To view this statement in the configuration.</p> <p>admin-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 99

traceoptions (Uplink Failure Detection)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <no-stamp> <replace> <size <i>size</i>> <world-readable no-world-readable>; flag <i>flag</i>; }</pre>
Hierarchy Level	[edit protocols uplink-failure-detection]
Release Information	Statement introduced in Junos OS Release 12.1 for EX Series switches.
Description	Define tracing operations for uplink failure detection.
Default	The traceoptions feature is disabled by default.
Options	<p>file <i>filename</i> —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks.</p> <p>files <i>number</i> —(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached (xk to specify KB, xm to specify MB, or xg to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the size option.</p> <p>Range: 2 through 1000</p> <p>Default: 3 files</p> <p>flag <i>flag</i> —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none">• all—Trace everything.• dcd—Trace ufdi interaction with dcd.• groups—Trace uplink-failure-detection group handling.• interface—Trace interface notification handlers of ufdi.• parse—Trace configuration parsing. <p>no-stamp—(Optional) Do not place a timestamp on any trace file.</p> <p>no-world-readable—(Optional) Restricted file access to the user who created the file.</p> <p>size <i>size</i> —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the files option.</p> <p>Syntax: xk to specify KB, xm to specify MB, or xg to specify gigabytes</p>

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

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

Related Documentation • [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 99](#)

uplink-failure-detection

```
Syntax  uplink-failure-detection {
        action {
            log;
        }
        group group-name {
            link-to-monitor interface-name;
            link-to-disable interface-name;
        }
        traceoptions {
            file filename <files number> <no-stamp> <replace> <size size> <world-readable |
            no-world-readable>;
            flag flag);
        }
    }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description Configure uplink and downlink interfaces in a group to monitor uplink failures and to propagate uplink failures to the downlink interfaces.

The remaining statements are explained separately.

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

Related Documentation • [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 99](#)

PART 3

Administration

- [Routine Monitoring on page 297](#)
- [Operational Commands: General on page 305](#)
- [Operational Commands: RPM on page 329](#)
- [Operational Commands: SNMP on page 341](#)
- [Operational Commands: Port Mirroring on page 379](#)
- [Operational Commands: sFlow on page 383](#)
- [Operational Commands: Ethernet OAM Connectivity Fault Management on page 391](#)
- [Operational Commands: Ethernet OAM Link Fault Management on page 431](#)
- [Operational Commands: Uplink Failure Detection on page 437](#)

Routine Monitoring

- [Monitoring Hosts Using the J-Web Ping Host Tool on page 297](#)
- [Monitoring Network Traffic Using Traceroute on page 299](#)
- [Verifying Input and Output for Port Mirroring Analyzers on EX Series Switches on page 300](#)
- [Viewing Real-Time Performance Monitoring Information on page 302](#)
- [Verifying That Uplink Failure Detection Is Working Correctly on page 302](#)

Monitoring Hosts Using the J-Web Ping Host Tool

Purpose



NOTE: This topic applies only to the J-Web Application package.

Use the J-Web ping host tool to verify that the host can be reached over the network. The output is useful for diagnosing host and network connectivity problems. The switch sends a series of ICMP echo (ping) requests to a specified host and receives ICMP echo responses.

Action

To use the J-Web ping host tool:

1. Select **Troubleshoot>Ping Host**.
2. Next to Advanced options, click the expand icon.
3. Enter information into the Ping Host page, as described in [Table 10 on page 298](#).

The Remote Host field is the only required field.

4. Click **Start**.

The results of the ping operation are displayed 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. To stop the ping operation before it is complete, click **OK**.

Meaning

[Table 10 on page 298](#) lists the fields.

Table 10: J-Web Ping Host Field Summary

Field	Function	Your Action
Remote Host	Identifies the host to ping.	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"> To suppress the display of the hop hostnames, select the check box. To display the hop hostnames, clear 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"> To set the DF bit, select the check box. To clear the DF bit, clear the check box.
Record Route	Sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed in the main pane.	<ul style="list-style-type: none"> To record and display the path of the packet, select the check box. To suppress the recording and display of the path of the packet, clear 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	Name of the routing instance for the ping attempt.	Select the routing instance name from the list.
Interval	Specifies the interval, in seconds, between transmissions of individual ping requests.	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 65468. The switch 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 value 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"> To bypass the routing table and send the ping requests to hosts on the specified interface only, select the check box. To route the ping requests using the routing table, clear the check box.

Related Documentation • *Monitoring Interface Status and Traffic*

Monitoring Network Traffic Using Traceroute

Purpose



NOTE: This topic applies only to the J-Web Application package.

Use the Traceroute page in the J-Web interface to trace a route between the switch and a remote host. You can use a traceroute task to display a list of waypoints between the switch and a specified destination host. The output is useful for diagnosing a point of failure in the path from the switch platform to the destination host and addressing network traffic latency and throughput problems.

Action

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.
The **Remote Host** field is the only required field.
4. Click **Start**.
5. To stop the traceroute operation before it is complete, click **OK** while the results of the traceroute operation are being displayed.

Meaning

The switch generates the list of waypoints by sending a series of ICMP traceroute packets in which the time-to-live (TTL) value in the messages sent to each successive waypoint is incremented by 1. (The TTL value of the first traceroute packet is set to 1.) In this manner, each waypoint along the path to the destination host replies with a Time Exceeded packet from which the source IP address can be obtained.

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 switch sends a total of three traceroute packets to each waypoint along the path and displays the round-trip time for each traceroute operation. If the switch times out before receiving a **Time Exceeded** message, an asterisk (*) is displayed for that round-trip time.

Table 11: Traceroute field summary

Field	Function	Your Action
Remote Host	Identifies the destination host of the traceroute.	Type the hostname or IP address of the destination host.

Table 11: Traceroute field summary (*continued*)

Field	Function	Your Action
Advanced Options		
Don't Resolve Addresses	Determines whether hostnames of the hops along the path are displayed, in addition to IP addresses.	To suppress the display of the hop hostnames, select the check box.
Gateway	Specifies the IP address of the gateway to route through.	Type the gateway IP address.
Source Address	Specifies the source address of the outgoing traceroute packets.	Type the source IP address.
Bypass Routing	Determines whether traceroute packets are routed by means of the routing table. 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.	To bypass the routing table and send the traceroute packets to hosts on the specified interface only, select the check box.
Interface	Specifies the interface on which the traceroute packets are sent.	From the list, select the interface on which traceroute packets are sent. 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.	From the list, select the TTL.
Type-of-Service	Specifies the type-of-service (TOS) value to include in the IP header of the traceroute request packet.	From the list, select the decimal value of the TOS field.
Resolve AS Numbers	Determines whether the autonomous system (AS) number of each intermediate hop between the router and the destination host is displayed.	To display the AS numbers, select the check box.

- Related Documentation**
- [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#)
 - [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#)
 - [Configuring Gigabit Ethernet Interfaces \(J-Web Procedure\)](#)
 - [Monitoring Interface Status and Traffic](#)

Verifying Input and Output for Port Mirroring Analyzers on EX Series Switches

- Purpose** Verify that an analyzer has been created on the switch and has the appropriate output interfaces, and appropriate output interface.

Action You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
[edit]
user@switch> show analyzer
Analyzer name           : employee-monitor
Output VLAN             : remote-analyzer
Mirror ratio            : 1
Loss priority           : High
Ingress monitored interfaces : ge-0/0/0.0
Ingress monitored interfaces : ge-0/0/1.0
```

You can view all of the port mirror analyzers configured on the switch, including any that are disabled, using the **show ethernet-switching-options** command in configuration mode.

```
user@switch# show ethernet-switching-options
inactive: analyzer employee-web-monitor {
    loss-priority high;
    output {

analyzer employee-monitor {
    loss-priority high;
    input {
        ingress {
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
        }
    }
    output {
        vlan {
            remote-analyzer;
        }
    }
}
```

Meaning This output shows that the employee-monitor analyzer has a ratio of 1 (mirroring every packet, the default), a loss priority of high (set this option to high whenever the analyzer output is to a VLAN), is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1**, and sending the mirrored traffic to the analyzer called remote-analyzer.

Related Documentation

- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
- [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)
- [Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on EX Series Switches on page 37](#)
- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
- [Understanding Port Mirroring on EX Series Switches on page 9](#)

Viewing Real-Time Performance Monitoring Information



NOTE: This topic applies only to the J-Web Application package.

Real-time performance monitoring (RPM) on EX Series switches enables you to configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter. The J-Web interface provides a graphical view of RPM information for EX Series switches.

To view the RPM information using the J-Web interface:

1. Select **Troubleshoot > RPM > View RPM**.
2. Select the **Round Trip Time** check box to display the graph with round-trip time included. Clear the check-box to view the graph without the round-trip time.
3. From the **Refresh Time** list, select a refresh time interval for the graph.

**Related
Documentation**

- [Configuring Real-Time Performance Monitoring \(J-Web Procedure\) on page 85](#)

Verifying That Uplink Failure Detection Is Working Correctly

Purpose Verify that the switch disables the downlink interface when it detects an uplink failure.

- Action** 1. View the current uplink-failure-detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : ge-0/0/0*
Downlink             : ge-0/0/1*
Failure Action       : Inactive
```



NOTE: The asterisk (*) indicates that the link is up.

2. Disable the uplink interface:

```
[edit]
user@switch# set interface ge-0/0/0 disable
```

3. Save the configuration on the switch.

4. View the current uplink-failure-detection status:

```
user@switch> show uplink-failure-detection
Group                : group1
Uplink               : ge-0/0/0
Downlink             : ge-0/0/1
Failure Action       : Active
```

Meaning The output in Step 1 shows that the uplink interface is up, and hence that the downlink interface is also up, and that the status of **Failure Action** is **Inactive**.

The output in Step 4 shows that both the uplink and downlink interfaces are down and that the status of **Failure Action** is changed to **Active**. This output shows that uplink failure detection is working.

- Related Documentation**
- [Configuring Interfaces for Uplink Failure Detection \(CLI Procedure\) on page 99](#)
 - [Understanding Uplink Failure Detection on page 27](#)

CHAPTER 16

Operational Commands: General

- monitor traffic
- ping
- show pfe statistics bridge
- traceroute

monitor traffic

Syntax monitor traffic
 <brief | detail | extensive>
 <absolute-sequence>
 <count *count*>
 <interface *interface-name*>
 <layer2-headers>
 <matching *matching*>
 <no-domain-names>
 <no-promiscuous>
 <no-resolve>
 <no-timestamp>
 <print-ascii>
 <print-hex>
 <resolve-timeout>
 <size *size*>

Release Information Command introduced before Junos OS Release 7.4.
 Command introduced in Junos OS Release 9.0 for EX Series switches.
 Command introduced in Junos OS Release 11.1 for the QFX Series.

Description Display packet headers or packets received and sent from the Routing Engine.



NOTE:

- Using the **monitor-traffic** command can degrade router or switch performance.
- Delays from DNS resolution can be eliminated by using the **no-resolve** option.



NOTE: This command is not supported on the QFabric system.

Options **none**—(Optional) Display packet headers transmitted through **fxp0**. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

brief | detail | extensive—(Optional) Display the specified level of output.

absolute-sequence—(Optional) Display absolute TCP sequence numbers.

count *count*—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The monitor traffic command quits automatically after displaying the number of packets specified.

interface *interface-name*—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

layer2-headers—(Optional) Display the link-level header on each line.

matching *matching*—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

no-domain-names—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only **team** for the hostname **team.company.net**.

no-promiscuous—(Optional) Do not put the interface into promiscuous mode.

no-resolve—(Optional) Suppress reverse lookup of the IP addresses.

no-timestamp—(Optional) Suppress timestamps on displayed packets.

print-ascii—(Optional) Display each packet in ASCII format.

print-hex—(Optional) Display each packet, except the link-level header, in hexadecimal format.

resolve-timeout *timeout*—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for 1 through 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

size *size*—(Optional) Read but do not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

Additional Information In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

```
monitor traffic matching "expression"
```

Replace ***expression*** with one or more of the match conditions listed in [Table 12 on page 308](#).

Table 12: Match Conditions for the monitor traffic Command

Match Type	Condition	Description
Entity	host [<i>address</i> <i>hostname</i>]	Matches packets that contain the specified address or hostname. The protocol match conditions arp , ip , or rarp , or any of the directional match conditions can be prepended to the host match condition.
	net <i>address</i>	Matches packets with source or destination addresses containing the specified network address.
	net <i>address</i> mask <i>mask</i>	Matches packets containing the specified network address and subnet mask.
	port (<i>port-number</i> <i>port-name</i>)	Matches packets containing the specified source or destination TCP or UDP port number or port name. In place of the numeric port address, you can specify a text synonym, such as bgp (179), dhcp (67), or domain (53) (the port numbers are also listed).
Directional	dst	Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.
	src	Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.
	src and dst	Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.
	src or dst	Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.
Packet Length	less <i>value</i>	Matches packets shorter than or equal to the specified value, in bytes.
	greater <i>value</i>	Matches packets longer than or equal to the specified value, in bytes.

Table 12: Match Conditions for the monitor traffic Command (*continued*)

Match Type	Condition	Description
Protocol	amt	Matches all AMT packets. Use the extensive level of output to decode the inner IGMP packets in addition to the AMT outer packet.
	arp	Matches all ARP packets.
	ether	Matches all Ethernet packets.
	ether (broadcast multicast)	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with src and dst .
	ether protocol (address (arp ip rarp))	Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The ether protocol arguments arp , ip , and rarp are also independent match conditions, so they must be preceded by a backslash (\) when used in the ether protocol match condition.
	icmp	Matches all ICMP packets.
	ip	Matches all IP packets.
	ip (broadcast multicast)	Matches broadcast or multicast IP packets.
	ip protocol (address (icmp igmp tcp udp))	Matches packets with the specified address or protocol type. The ip protocol arguments icmp , tcp , and udp are also independent match conditions, so they must be preceded by 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 datagrams.
	udp	Matches all UDP datagrams.

To combine expressions, use the logical operators listed in [Table 13 on page 309](#).

Table 13: Logical Operators for the monitor traffic Command

Logical Operator (Highest to Lowest Precedence)	Description
!	Logical NOT. If the first condition does not match, the next condition is evaluated.

Table 13: Logical Operators for the monitor traffic Command (*continued*)

Logical Operator (Highest to Lowest Precedence)	Description
&&	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 (\).

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

```
monitor traffic matching "ether[0] & 1 != 0"arithmetic_expression relational_operator arithmetic_expression
```

The packet data accessor uses the following syntax:

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

The optional *size* field represents the number of bytes examined in the packet header. The available values are 1, 2, or 4 bytes. The following sample command captures all multicast traffic:

```
user@host> monitor traffic matching "ether[0] & 1 != 0"
```

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in [Table 14 on page 311](#).



NOTE: Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The **monitor traffic** command cannot apply match conditions to inbound traffic.
- The **monitor traffic interface** command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (`| match`) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see [Table 12 on page 308](#).
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. Because the **monitor traffic interface ae[x]** command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

Table 14: Arithmetic and Relational Operators for the monitor traffic Command

Arithmetic or Relational Operator	Description
Arithmetic Operator	
+	Addition operator.
-	Subtraction operator.
/	Division operator.
&	Bitwise AND.
*	Bitwise exclusive OR.
	Bitwise inclusive OR.
Relational Operator (Highest to Lowest Precedence)	
<=	If the first expression is less than or equal to the second, the packet matches.
>=	If the first expression is greater than or equal to the second, the packet matches.
<	If the first expression is less than the second, the packet matches.
>	If the first expression is greater than the second, the packet matches.
=	If the compared expressions are equal, the packet matches.
!=	If the compared expressions are unequal, the packet matches.

Required Privilege Level trace
maintenance

List of Sample Output [monitor traffic count on page 312](#)
[monitor traffic detail count on page 312](#)
[monitor traffic extensive \(Absolute Sequence\) on page 312](#)
[monitor traffic extensive \(Relative Sequence\) on page 312](#)
[monitor traffic extensive count on page 312](#)
[monitor traffic interface on page 313](#)
[monitor traffic matching on page 313](#)
[monitor traffic \(TX Matrix Plus Router\) on page 313](#)
[monitor traffic \(QFX3500 Switch\) on page 314](#)

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185
Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864 In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926
Out my-server.work.net.telnet.telnet > my-server.home.net.1295: P 1:38(37) ack 0
win 17680 (DF) [tos 0x10] (ttl 6)
```

monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0] (ttl )
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53)
ack 1845421797 win 16384
<nop,nop,timestamp 4935628 965951>:
BGP [|BGP UPDAT)
In 192.168.4.227.1024 > 207.17.136.193.179:
P 1845421797:1845421852(55) ack 4042780912 win 16384 <nop,nop,timestamp 965951
4935628>: BGP [|BGP UPDAT)
...
```

monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0)
ack 1664980689 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10] (ttl 64, id 5394)
In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57)
ack 1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

monitor traffic extensive count

```
user@host> monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933
In 192.168.4.206.2723610880 > 172.17.28.8.2049:
40 null (ttl 64, id 38367)13:18:17.407577
In 172.17.28.8.2049 > 192.168.4.206.2723610880:
```

```

reply ok 28 null (ttl 61, id 35495)13:18:17.541140
In 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:
0000 0100 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 0000 0000
0000 0000 000013:18:17.591513
In 172.24.248.156.4139 > 192.168.4.210.23:
3556964918:3556964918(0)
ack 295526518 win 17601 (DF)
(ttl 121, id 14)13:18:17.591568
Out 192.168.4.210.23 >
172.24.248.156.4139: P 1:40(39)
ack 0 win 17680 (DF) [tos 0x10]
(ttl 64, id 52376)

```

monitor traffic interface

```

user@host> monitor traffic interface fxp0
listening on fxp0.0
18:17:28.800650 In server.home.net.723 > host1-0.lab.home.net.log
18:17:28.800733 Out host2-0.lab.home.net.login > server.home.net.7
18:17:28.817813 In host30.lab.home.net.syslog > host40.home0
18:17:28.817846 In host30.lab.home.net.syslog > host40.home0
...

```

monitor traffic matching

```

user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511 In IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing!
192.168.1.17.netbios-ns > 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864 In arp who-has 192.168.1.17 tell 192.168.1.9
...

```

monitor traffic (TX Matrix Plus Router)

```

user@host> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on em0, capture size 96 bytes
04:11:59.862121 Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog > sv-log-01.englab.juniper.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.englab.juniper.net.syslog >
sv-log-02.englab.juniper.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.englab.juniper.net.65235 >

```

```

summit-em0.englab.juniper.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.englab.juniper.net.telnet > aj-em0.englab.juniper.net.65235: P
1:241(240) ack 0 win 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.englab.juniper.net.65235 >
summit-em0.englab.juniper.net.telnet: . ack 241 win 33304 <nop,nop,timestamp
42366834 993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell11.juniper.net.46182 > summit-em0.englab.juniper.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell11.juniper.net.46182 >
summit-em0.englab.juniper.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.englab.juniper.net.telnet >
ipg-lnx-shell11.juniper.net.46182: P 28:37(9) ack 48 win 33304 <nop,nop,timestamp
993810 1308555294>
...

```

monitor traffic (QFX3500 Switch)

```

user@switch> monitor traffic
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.

```



```
Listening on me4, capture size 96 bytes
Reverse lookup for 172.22.16.246 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use <no-resolve> to avoid reverse lookups on IP addresses.
16:35:32.240873 Out IP truncated-ip - 112 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 4200727624:4200727756(132) ack 2889954831 win 65535
16:35:32.240900 Out IP truncated-ip - 176 bytes missing!
labqfx-me0.lab4.juniper.net.ssh >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```


ping

List of Syntax [Syntax on page 316](#)
 [Syntax \(QFX Series\) on page 316](#)

Syntax `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet | inet6>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict >`
 `<strict-source value.>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`
 `<vpls instance-name>`
 `<wait seconds>`

Syntax (QFX Series) `ping host`
 `<bypass-routing>`
 `<count requests>`
 `<detail>`
 `<do-not-fragment>`
 `<inet>`
 `<interface source-interface>`
 `<interval seconds>`
 `<logical-system logical-system-name>`
 `<loose-source value>`
 `<mac-address mac-address>`
 `<no-resolve>`
 `<pattern string>`
 `<rapid>`
 `<record-route>`
 `<routing-instance routing-instance-name>`
 `<size bytes>`
 `<source source-address>`
 `<strict>`
 `< strict-source value>`
 `<tos type-of-service>`
 `<ttl value>`
 `<verbose>`

<wait *seconds*>

Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Check host reachability and network connectivity. The ping command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.</p>
Options	<p>host—IP address or hostname of the remote system to ping.</p> <p>bypass-routing—(Optional) Bypass the normal routing tables and send ping 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 ping a local system through an interface that has no route through it.</p> <p>count requests—(Optional) Number of ping requests to send. The range of values is 1 through 2,000,000,000. The default value is an unlimited number of requests.</p> <p>detail—(Optional) Include in the output the interface on which the ping reply was received.</p> <p>do-not-fragment—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets. For IPv6 packets, this option disables fragmentation.</p> <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <p> NOTE: In Junos OS Release 11.1 and later, when issuing the ping command for an IPv6 route with the do-not-fragment option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.</p> </div> <p>inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.</p> <p>inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.</p> <p>interface source-interface—(Optional) Interface to use to send the ping requests.</p> <p>interval seconds—(Optional) How often to send ping requests. The range of values, in seconds, is 1 through infinity. The default value is 1.</p> <p>logical-system logical-system-name—(Optional) Name of logical system from which to send the ping requests.</p> <p>Alternatively, enter the set cli logical-system logical-system-name command and then run the ping command. To return to the main router or switch, enter the clear cli logical-system command.</p>

loose-source value—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

mac-address mac-address—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern string—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid—(Optional) Send 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) Record and report the packet's path (IPv4).

routing-instance routing-instance-name—(Optional) Name of the routing instance for the ping attempt.

size bytes—(Optional) Size of ping request packets. The range of values, in bytes, is 0 through 65,468. The default value is 56, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

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).

strict—(Optional) Use the strict source route option (IPv4).

strict-source value—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos type-of-service—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is 0 through 255.

If the device configuration includes the **dscp-code-point value** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

ttl value—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is 0 through 255.

verbose—(Optional) Display detailed output.

vpls instance-name—(Optional) Ping the instance to which this VPLS belongs.

wait seconds—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is 10 seconds. If this option is used without the count option, a default count of 5 packets is used.

Required Privilege Level	network
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages</i>
List of Sample Output	ping hostname on page 319 ping hostname rapid on page 319 ping hostname size count on page 319
Output Fields	<p>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.</p>

Sample Output

ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

show pfe statistics bridge

Syntax	show pfe statistics bridge <fpc slot>
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display information about the number of packets discarded in the ingress pipeline of the Packet Forwarding Engine, packets discarded because of egress filtering or congestion filtering, number of control packets, and general counters for dropped packets. You can use this information to inform troubleshooting investigations.
Options	none —Display bridge counter statistics for all Flexible PIC Concentrator (FPC) slots. fpc slot —(Optional) Display bridge counter statistics for a specific FPC slot.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Monitoring System Process Information</i> • <i>Monitoring Switch Control Traffic</i>
List of Sample Output	show pfe statistics bridge (EX3200 and EX4200 Switches) on page 321 show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis) on page 322 show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis) on page 323 show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card) on page 323
Output Fields	Table 15 on page 320 lists the output fields for the show pfe statistics bridge command. Output fields are listed in the approximate order in which they appear.

Table 15: show pfe statistics bridge Output Fields

Field Name	Field Description
Ingress Counters	<p>Information about ingress counters:</p> <ul style="list-style-type: none"> • Received—Number of packets received by the bridge. • VLAN Filtered—Number of packets discarded because of VLAN filtering. • Security Filtered—Number of packets discarded because of security filtering. • Other Discards—Number of packets dropped by the bridge for reasons other than VLAN or security filtering.

Table 15: show pfe statistics bridge Output Fields (*continued*)

Field Name	Field Description
Egress Counters	Information about egress counters: <ul style="list-style-type: none"> • Unicast—Number of unicast packets transmitted. • Multicast—Number of multicast packets transmitted. • Broadcast—Number of broadcast packets transmitted. • Egress Filtered—Number of egress-filtered packets (regardless of port, priority, or mode). • TailDrop—Number of packets filtered because of egress queue congestion. • Forward Restrict—Number of packets filtered because of egress forward restrictions. • Congestion Filtered—Number of packets filtered because of transmit queue (TxQ) congestion. • Control Packets—Number of control packets (sent to CPU, received from CPU, and sent to analyzer).
Drop Counters	Information about drop counters: <ul style="list-style-type: none"> • Drop Mode—Count mode of the counter. • Drop Counter—Counter value.
General Counters	Information about general counters: <ul style="list-style-type: none"> • Drop Mode—Count mode of the counter. • Drop Counter—Counter value. • Source Not Learnt—Number of source addresses that were not learnt because of internal congestion.
MUX PFE	Information about multiplexer PFE for oversubscribed cards: <ul style="list-style-type: none"> • Drop Mode—Count mode of the counter. • Drop Count—Counter value.

Sample Output

show pfe statistics bridge (EX3200 and EX4200 Switches)

```

user@switch> show pfe statistics bridge
Slot 0

PFE:                0          1          2
-----
---- Ingress Counters ----
Received:            0          52          0
VLAN Filtered:       0          0          0
Security Filtered:   0          0          0
Other Discards:      0          0          0
---- Egress Counters ----
Unicast:             0         104         52
Multicast:           0          0          0
Broadcast:           0          0          0
Egress Filtered:     0          0          0
Congestion Filtered: 0          0          0
Control Packets:     5          0          0
---- General Counters ----
Drop Mode:           0          0          0

```

```

Drop Counter:          34217      36080      6367
Source Not Learnt:      0          0          0

```

show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis)

```
user@switch> show pfe statistics bridge
```

```
Slot 0
```

```
PFE:                  0          1
```

```
----- Ingress Counters -----
```

```

Received:              946          48
VLAN Filtered:         0           0
Security Filtered:     0           0
Other Discards:        0           0

```

```
----- Egress Counters -----
```

```

Unicast:               0           0
Multicast:             0           0
Broadcast:             0           0
Egress Filtered:       0           0
TailDrop:              0           0
Forward Restrict:      0           0
Congestion Filtered:   0           0
Control Packets:       4103        896

```

```
----- Drop Counters -----
```

```

Drop Mode:             0           0
Drop Counter:          12528        2

```

```
Slot 1
```

```
PFE:                  0          1
```

```
----- Ingress Counters -----
```

```

Received:              0           0
VLAN Filtered:         0           0
Security Filtered:     0           0
Other Discards:        0           0

```

```
----- Egress Counters -----
```

```

Unicast:               0           0
Multicast:             0           0
Broadcast:             0           0
Egress Filtered:       0           0
TailDrop:              0           0
Forward Restrict:      0           0
Congestion Filtered:   0           0
Control Packets:       0           0

```

```
----- Drop Counters -----
```

```

Drop Mode:             0           0
Drop Counter:          0           0

```

```
Slot 2
```

```
PFE:                  0          1
```

```
----- Ingress Counters -----
```

```

Received:              0           0
VLAN Filtered:         0           0
Security Filtered:     0           0
Other Filtered:        0           0

```

```
----- Egress Counters -----
```

```

Unicast:               0           0
Multicast:             0           0

```



```

Broadcast:                0          0
Egress Filtered:          0          0
TailDrop:                 0          0
Forward Restrict:         0          0
Congestion Filtered:      0          0
Control Packets:          0          0
---- Drop Counters ----
Drop Mode:                0          0
Drop Counter:             0          0

```

show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis)

```

user@switch> show pfe statistics bridge fpc 2
Slot 2

```

```

PFE:                0          1
-----
---- Ingress Counters ----
Received:           0          0
VLAN Filtered:      0          0
Security Filtered:  0          0
Other Discards:     0          0
---- Egress Counters ----
Unicast:            0          0
Multicast:          0          0
Broadcast:          0          0
Egress Filtered:    0          0
TailDrop:           0          0
Forward Restrict:   0          0
Congestion Filtered: 0          0
Control Packets:    0          0
---- Drop Counters ----
Drop Mode:          0          0
Drop Counter:       0          0

```

show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card)

```

user@switch> show pfe statistics bridge fpc 8
Slot 8

```

```

PFE:                0          1          2          3
-----
---- Ingress Counters ----
Received:           0          3          0          0
VLAN Filtered:      0          0          0          0
Security Filtered:  0          0          0          0
Other Discards:     0          1          0          0
---- Egress Counters ----
Unicast:            0          0          0          0
Multicast:          0          0          0          0
Broadcast:          0          0          0          0
Egress Filtered:    0          0          0          0
TailDrop:           0          0          0          0
Forward Restrict:   0          0          0          0
Congestion Filtered: 0          2          0          0
Control Packets:    4          0          0          0
---- Drop Counters ----
Drop Mode:          0          0          0          0
Drop Counter:       0          1          0          0

MUX PFE:            4          5

```

Drop Mode: 0 0
Drop Count: 0 0

traceroute

List of Syntax [Syntax on page 325](#)
 [Syntax \(QFX Series\) on page 325](#)

Syntax `traceroute host`
 `<as-number-lookup>`
 `<bypass-routing>`
 `<clns>`
 `<gateway address>`
 `<inet | inet6>`
 `<interface interface-name>`
 `<logical system logical-system-name>`
 `<monitor host>`
 `<mpls (ldp FEC address | rsvp label-switched-path-name)>`
 `<no-resolve>`
 `<propagate-ttl>`
 `<routing-instance routing-instance-name>`
 `<source source-address>`
 `<tos value>`
 `<ttl value>`
 `<wait seconds>`

Syntax (QFX Series) `traceroute host`
 `<as-number-lookup>`
 `<bypass-routing>`
 `<gateway address>`
 `<inet>`
 `<interface interface-name>`
 `<monitor host>`
 `<no-resolve>`
 `<routing-instance routing-instance-name>`
 `<source source-address>`
 `<tos value>`
 `<ttl value>`
 `<wait seconds>`

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.

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 327](#)
- [traceroute as-number-lookup host on page 327](#)
- [traceroute no-resolve on page 327](#)
- [traceroute propagate-ttl on page 328](#)
- [traceroute \(Between CE Routers, Layer 3 VPN\) on page 328](#)
- [traceroute \(Through an MPLS LSP\) on page 328](#)

Output Fields [Table 16 on page 327](#) describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 16: 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 green.company.net (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 green.company.net (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 vpn09.skybank.net (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  vpn09.skybank.net (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
```

CHAPTER 17

Operational Commands: RPM

- `show services rpm active-servers`
- `show services rpm history-results`
- `show services rpm probe-results`

show services rpm active-servers

Syntax	show services rpm active-servers
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 13.2 for PTX Series Packet Transport Routers.
Description	Display the protocols and corresponding ports for which a router or switch is configured as a real-time performance monitoring (RPM) server.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show services rpm active-servers on page 330
Output Fields	Table 17 on page 330 lists the output fields for the show services rpm active-servers command. Output fields are listed in the approximate order in which they appear.

Table 17: show services rpm active-servers Output Fields

Field Name	Field Description
Protocol	Protocol configured on the receiving probe server. The protocol can be the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP).
Port	Port configured on the receiving probe server.
Destination interface name	Output interface name for the probes.

Sample Output

show services rpm active-servers

```
user@host> show services rpm active-servers
  Protocol: TCP, Port: 50000, Destination interface name: lt-0/0/0.0
  Protocol: UDP, Port: 50001, Destination interface name: lt-0/0/0.0
```


show services rpm history-results

Syntax	<pre>show services rpm history-results <brief detail> <owner <i>owner</i>> <since <i>time</i>> <test <i>name</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.</p>
Description	Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) instance.
Options	<p>none—Display the results of the last 50 probes for all RPM instances.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>owner <i>owner</i>—(Optional) Display information for the specified probe owner.</p> <p>since <i>time</i>—(Optional) Display information from the specified time. Specify time as <i>yyyy-mm-dd.hh:mm:ss</i>.</p> <p>test <i>name</i>—(Optional) Display information for the specified test.</p>
Required Privilege Level	view
List of Sample Output	<p>show services rpm history-results on page 332</p> <p>show services rpm history-results detail on page 332</p>
Output Fields	Table 18 on page 331 lists the output fields for the show services rpm history-results command. Output fields are listed in the approximate order in which they appear.

Table 18: show services rpm history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Probe owner.	All levels
Test	Name of a test for a probe instance.	All levels
Probe received	Timestamp when the probe result was determined.	All levels
Round trip time	Average ping round-trip time (RTT), in microseconds.	All levels
Probe results	<p>Result of a particular probe performed by a remote host. The following information is contained in the results:</p> <ul style="list-style-type: none"> Response received—Timestamp when the probe result was determined. Rtt—Average ping round-trip time (RTT), in microseconds. 	detail

Table 18: show services rpm history-results Output Fields (*continued*)

Field Name	Field Description	Level of Output
Results over current test	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	detail
Probes sent	Number of probes sent with the current test.	detail
Probes received	Number of probe responses received within the current test.	detail
Loss percentage	Percentage of lost probes for the current test.	detail
Measurement	<p>Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-pin-timestamp, the egress and ingress delay:</p> <ul style="list-style-type: none"> • 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. • Jitter—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • Stddev—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test. 	detail

Sample Output

show services rpm history-results

```

user@host> show services rpm history-results
      Owner, Test                Probe received                Round trip time
p1, t1                          Wed Aug 12 01:02:35 2009          315 usec
p1, t1                          Wed Aug 12 01:02:36 2009          266 usec
p1, t1                          Wed Aug 12 01:02:37 2009          314 usec
p1, t1                          Wed Aug 12 01:02:38 2009          388 usec
p1, t1                          Wed Aug 12 01:02:39 2009          316 usec
p1, t1                          Wed Aug 12 01:02:40 2009          271 usec
p1, t1                          Wed Aug 12 01:02:41 2009          314 usec
p1, t1                          Wed Aug 12 01:02:42 2009          1180 usec

```

show services rpm history-results detail

```

user@host> show services rpm history-results detail
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
  Response received, Wed Aug 12 01:02:35 2009,
  Client and server hardware timestamps
  Rtt: 315 usec
Results over current test:
  Probes sent: 1, Probes received: 1, Loss percentage: 0
  Measurement: Round trip time
    Samples: 1, Minimum: 315 usec, Maximum: 315 usec, Average: 315 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 315 usec

```

```
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
  Response received, Wed Aug 12 01:02:36 2009,
  Client and server hardware timestamps
  Rtt: 266 usec, Round trip jitter: -50 usec,
  Round trip interarrival jitter: 3 usec
Results over current test:
  Probes sent: 2, Probes received: 2, Loss percentage: 0
  Measurement: Round trip time
    Samples: 2, Minimum: 266 usec, Maximum: 315 usec, Average: 291 usec,
    Peak to peak: 49 usec, Stddev: 24 usec, Sum: 581 usec
  Measurement: Negative round trip jitter
    Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
  Response received, Wed Aug 12 01:02:37 2009,
  Client and server hardware timestamps
  Rtt: 314 usec, Round trip jitter: 49 usec,
  Round trip interarrival jitter: 6 usec
Results over current test:
  Probes sent: 3, Probes received: 3, Loss percentage: 0
  Measurement: Round trip time
    Samples: 3, Minimum: 266 usec, Maximum: 315 usec, Average: 298 usec,
    Peak to peak: 49 usec, Stddev: 23 usec, Sum: 895 usec
  Measurement: Positive round trip jitter
    Samples: 1, Minimum: 49 usec, Maximum: 49 usec, Average: 49 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 49 usec
  Measurement: Negative round trip jitter
    Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
  Response received, Wed Aug 12 01:02:38 2009,
  Client and server hardware timestamps
  Rtt: 388 usec, Round trip jitter: 74 usec,
  Round trip interarrival jitter: 10 usec
Results over current test:
  Probes sent: 4, Probes received: 4, Loss percentage: 0
  Measurement: Round trip time
    Samples: 4, Minimum: 266 usec, Maximum: 388 usec, Average: 321 usec,
    Peak to peak: 122 usec, Stddev: 44 usec, Sum: 1283 usec
  Measurement: Positive round trip jitter
    Samples: 2, Minimum: 49 usec, Maximum: 74 usec, Average: 62 usec,
    Peak to peak: 25 usec, Stddev: 12 usec, Sum: 123 usec
  Measurement: Negative round trip jitter
    Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec
```

show services rpm probe-results

Syntax	show services rpm probe-results <owner <i>owner</i> > <test <i>name</i> >
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 13.2 for PTX Series Packet Transport Series Routers.
Description	Display the results of the most recent real-time performance monitoring (RPM) probes.
Options	none —Display all results of the most recent RPM probes. owner <i>owner</i> —(Optional) Display information for the specified probe owner. test <i>name</i> —(Optional) Display information for the specified test.
Required Privilege Level	view
List of Sample Output	show services rpm probe-results on page 337 show services rpm probe-results (BGP Neighbor Discovery) on page 339
Output Fields	Table 19 on page 334 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 19: 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 address used for the probes.
Source address	Source address used for the probes.
Probe type	Protocol configured on the receiving probe server: http-get , http-metadata-get , icmp-ping , icmp-ping-timestamp , tcp-ping , udp-ping , or udp-ping-timestamp .
Test size	Number of probes within a test.

Table 19: show services rpm probe-results Output Fields (*continued*)

Field Name	Field Description
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.
Results over current test	<p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent within the current test. Probes received—Number of probe responses received within the current test. Loss percentage—Percentage of lost probes for the current test. Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. Peak to peak—Peak-to-peak difference, in microseconds. Stddev—Standard deviation, in microseconds. Sum—Statistical sum.

Table 19: show services rpm probe-results Output Fields (*continued*)

Field Name	Field Description
Results over last test	<p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent for the most recently completed test. • Probes received—Number of probe responses received for the most recently completed test. • Loss percentage—Percentage of lost probes for the most recently completed test. • Test completed—Time the most recent test was completed. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following 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.
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.

Table 19: show services rpm probe-results Output Fields (*continued*)

Field Name	Field Description
Error Stats	<p>Displays error statistics for each probe.</p> <ul style="list-style-type: none"> • Invalid client rcv timestamp—Number of client receive timestamp less than client send timestamp. • Invalid server send timestamp—Number of server send timestamp less than server receive timestamp. • Invalid server processing time—Number of server side spent time greater than RTT. <p>NOTE: Error Stats is displayed in the output only if non-zero statistics exists.</p>

Sample Output

show services rpm probe-results

```

user@host> show services rpm probe-results
Owner: ADSN-J4300.ADSN-J2300.D2, Test: 75300002
Target address: 172.16.54.172, Source address: 10.206.0.1,
Probe type: udp-ping-timestamp, Test size: 10 probes
Probe results:
  Response received, Tue Feb  6 14:53:15 2007,
  Client and server hardware timestamps
  Rtt: 575 usec, Egress jitter: 5 usec, Ingress jitter: 8 usec,
  Round trip jitter: 12 usec, Egress interarrival jitter: 8 usec,
  Ingress interarrival jitter: 7 usec, Round trip interarrival jitter: 7 usec,

  Round trip interarrival jitter: 669 usec
Results over current test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0
  Measurement: Round trip time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive round trip jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative round trip jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Egress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over last test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0

```

```
Test completed on Tue Feb 6 14:53:16 2007
Measurement: Round trip time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Egress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over all tests:
Probes sent: 560, Probes received: 560, Loss percentage: 0
Measurement: Round trip time
  Samples: 560, Minimum: 805 usec, Maximum: 3114 usec, Average: 1756 usec,

  Peak to peak: 2309 usec, Stddev: 519 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
  Samples: 257, Minimum: 0 usec, Maximum: 2054 usec, Average: 597 usec,
  Peak to peak: 2054 usec, Stddev: 427 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
  Samples: 302, Minimum: 1 usec, Maximum: 1812 usec, Average: 511 usec,
  Peak to peak: 1811 usec, Stddev: 408 usec, Sum: xxxx usec
Measurement: Egress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
  Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
  Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
  Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
  Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
  Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Error Stats:
  Invalid client rcv timestamp: 3, Invalid server send timestamp: 0
  Invalid server processing time: 0
```


show services rpm probe-results (BGP Neighbor Discovery)

```
user@host> show services rpm probe-results
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: LS1/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
```


CHAPTER 18

Operational Commands: SNMP

- `clear snmp rmon history`
- `clear snmp statistics`
- `request snmp spoof-trap`
- `show snmp health-monitor`
- `show snmp inform-statistics`
- `show snmp mib`
- `show snmp rmon`
- `show snmp rmon history`
- `show snmp statistics`
- `show snmp v3`

clear snmp rmon history

Syntax	clear snmp rmon history < <i>interface-name</i> all>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	<p>Delete the samples of Ethernet statistics collected, but do not delete the RMON history configuration.</p> <p>The clear snmp rmon history command deletes all the samples collected for the interface configured for the history group, but not the configuration of that group. If you want to delete the RMON history group configuration, you must use the delete snmp rmon history configuration-mode command.</p>
Options	<p>interface-name—Delete the samples of Ethernet statistics collected for this interface.</p> <p>all—Delete the samples of Ethernet statistics collected for all interfaces that have been configured for RMON monitoring.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show snmp rmon history on page 367

clear snmp statistics

Syntax	clear snmp statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Clear Simple Network Management Protocol (SNMP) statistics.
Options	This command has no options.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show snmp statistics on page 371
List of Sample Output	clear snmp statistics on page 343
Output Fields	See show snmp statistics for an explanation of output fields.

Sample Output

clear snmp statistics

In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290
```

```
user@host> clear snmp statistics
```

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
```

```
Total request varbinds: 0, Total set varbinds: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0,  
Silent drops: 0, Proxy drops 0  
Output:  
Packets: 0, Too bigs: 0, No such names: 0,  
Bad values: 0, General errors: 0,  
Get requests: 0, Get nexts: 0, Set requests: 0,  
Get responses: 0, Traps: 0
```

request snmp spoof-trap

Syntax	request snmp spoof-trap <trap> variable-bindings <object> <instance> <value>
Release Information	Command introduced in Junos OS Release 8.2. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.
Options	<p><trap>—Name of the trap to spoof.</p> <p>variable-bindings <object> <instance> <value>—(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, ifIndex[14] = 14). Enclose the list of variable bindings in quotation marks (" ") and use a comma to separate each object name, instance, and value definition (for example, variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.</p> <p><dummy name>—A dummy trap name to display the list of available traps.</p> <p>Question mark (?)—Question mark? to display possible completions.</p>
Required Privilege Level	request
List of Sample Output	request snmp spoof-trap (with Variable Bindings) on page 345 request snmp spoof-trap (Illegal Trap Name) on page 345 request snmp spoof-trap (Question Mark ?) on page 349

Sample Output

request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmpspoof-trap linkUp variable-bindings "ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2"
Spoof trap request result: trap sent successfully
```

request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found
```

```
Allowed Traps:
ads1AtucInitFailureTrap
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
```

ads1AturPerfLofsThreshTrap
ads1AturPerfLossThreshTrap
ads1AturPerfLprsThreshTrap
ads1AturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlswTrapCircuitDown
dlswTrapCircuitUp
dlswTrapTConnDown
dlswTrapTConnPartnerReject
dlswTrapTConnProtViolation
dlswTrapTConnUp
dsx1LineStatusChange
dsx3LineStatusChange
entConfigChange
fallingAlarm
frDLCIStatusChange
ggsnTrapChanged
ggsnTrapCleared
ggsnTrapNew
gmp1sTunnelDown
ifMauJabberTrap
ipv6IfStateChange
isisAreaMismatch
isisAttemptToExceedMaxSequence
isisAuthenticationFailure
isisAuthenticationTypeFailure
isisCorruptedLSPDetected
isisDatabaseOverload
isisIDLenMismatch
isisLSPTooLargeToPropagate
isisManualAddressDrops
isisMaxAreaAddressesMismatch
isisOriginatingLSPBufferSizeMismatch
isisOwnLSPPurge
isisProtocolsSupportedMismatch
isisRejectedAdjacency
isisSequenceNumberSkip
isisVersionSkew
jnxAccessAuthServerDisabled
jnxAccessAuthServerEnabled
jnxAccessAuthServiceDown
jnxAccessAuthServiceUp
jnxBfdSessDetectionTimeHigh
jnxBfdSessTxIntervalHigh
jnxBgpM2BackwardTransition
jnxBgpM2Established
jnxCmCfgChange
jnxCmRescueChange
jnxCollFlowOverload
jnxCollFlowOverloadCleared
jnxCollFtpSwitchover

jnxCollMemoryAvailable
jnxCollMemoryUnavailable
jnxCollUnavailableDest
jnxCollUnavailableDestCleared
jnxCollUnsuccessfulTransfer
jnxDfcHardMemThresholdExceeded
jnxDfcHardMemUnderThreshold
jnxDfcHardPpsThresholdExceeded
jnxDfcHardPpsUnderThreshold
jnxDfcSoftMemThresholdExceeded
jnxDfcSoftMemUnderThreshold
jnxDfcSoftPpsThresholdExceeded
jnxDfcSoftPpsUnderThreshold
jnxEventTrap
jnxExampleStartup
jnxFEBSwitchover
jnxFanFailure
jnxFanOK
jnxFruCheck
jnxFruFailed
jnxFruInsertion
jnxFruOK
jnxFruOffline
jnxFruOnline
jnxFruPowerOff
jnxFruPowerOn
jnxFruRemoval
jnxHardDiskFailed
jnxHardDiskMissing
jnxJsAvPatternUpdateTrap
jnxJsChassisClusterSwitchover
jnxJsFwAuthCapacityExceeded
jnxJsFwAuthFailure
jnxJsFwAuthServiceDown
jnxJsFwAuthServiceUp
jnxJsNatAddrPoolThresholdStatus
jnxJsScreenAttack
jnxJsScreenCfgChange
jnxLdpLspDown
jnxLdpLspUp
jnxLdpSesDown
jnxLdpSesUp
jnxMIMstCistPortLoopProtectStateChangeTrap
jnxMIMstCistPortRootProtectStateChangeTrap
jnxMIMstErrTrap
jnxMIMstGenTrap
jnxMIMstInvalidBpduRxdTrap
jnxMIMstMstiPortLoopProtectStateChangeTrap
jnxMIMstMstiPortRootProtectStateChangeTrap
jnxMIMstNewRootTrap
jnxMIMstProtocolMigrationTrap
jnxMIMstRegionConfigChangeTrap
jnxMIMstTopologyChgTrap
jnxMacChangedNotification
jnxMplsLdpInitSesThresholdExceeded
jnxMplsLdpPathVectorLimitMismatch
jnxMplsLdpSessionDown
jnxMplsLdpSessionUp
jnxOspfV3IfConfigError
jnxOspfV3IfRxBadPacket
jnxOspfV3IfStateChange

jnxOspfV3LsdbApproachingOverflow
jnxOspfV3LsdbOverflow
jnxOspfV3NbrRestartHelperStatusChange
jnxOspfV3NbrStateChange
jnxOspfV3NssaTranslatorStatusChange
jnxOspfV3RestartStatusChange
jnxOspfV3VirtIfConfigError
jnxOspfV3VirtIfRxBadPacket
jnxOspfV3VirtIfStateChange
jnxOspfV3VirtNbrRestartHelperStatusChange
jnxOspfV3VirtNbrStateChange
jnxOtnAlarmCleared
jnxOtnAlarmSet
jnxOverTemperature
jnxPMonOverloadCleared
jnxPMonOverloadSet
jnxPingEgressJitterThresholdExceeded
jnxPingEgressStdDevThresholdExceeded
jnxPingEgressThresholdExceeded
jnxPingIngressJitterThresholdExceeded
jnxPingIngressStdDevThresholdExceeded
jnxPingIngressThresholdExceeded
jnxPingRttJitterThresholdExceeded
jnxPingRttStdDevThresholdExceeded
jnxPingRttThresholdExceeded
jnxPortBpduErrorStatusChangeTrap
jnxPortLoopProtectStateChangeTrap
jnxPortRootProtectStateChangeTrap
jnxPowerSupplyFailure
jnxPowerSupplyOK
jnxRedundancySwitchover
jnxRmonAlarmGetFailure
jnxRmonGetOk
jnxSecAccessIfMacLimitExceeded
jnxSecAccessSdsRateLimitCrossed
jnxSonetAlarmCleared
jnxSonetAlarmSet
jnxSpSvcSetCpuExceeded
jnxSpSvcSetCpuOk
jnxSpSvcSetZoneEntered
jnxSpSvcSetZoneExited
jnxStormEventNotification
jnxSyslogTrap
jnxTemperatureOK
jnxVccpPortDown
jnxVccpPortUp
jnxVpnIfDown
jnxVpnIfUp
jnxVpnPwDown
jnxVpnPwUp
jnxl2aldGlobalMacLimit
jnxl2aldInterfaceMacLimit
jnxl2aldRoutingInstMacLimit
linkDown
linkUp
lldpRemTablesChange
mfrMibTrapBundleLinkMismatch
mplsLspChange
mplsLspDown
mplsLspInfoChange
mplsLspInfoDown

```

mplsLspInfoPathDown
mplsLspInfoPathUp
mplsLspInfoUp
mplsLspPathDown
mplsLspPathUp
mplsLspUp
mplsNumVrfRouteMaxThreshExceeded
mplsNumVrfRouteMidThreshExceeded
mplsNumVrfSecIllglLb1ThrshExcd
mplsTunnelDown
mplsTunnelReoptimized
mplsTunnelRerouted
mplsTunnelUp
mplsVrfIfDown
mplsVrfIfUp
mplsXCDown
mplsXCUp
msdpBackwardTransition
msdpEstablished
newRoot
ospfIfAuthFailure
ospfIfConfigError
ospfIfRxBadPacket
ospfIfStateChange
ospfLsdbApproachingOverflow
ospfLsdbOverflow
ospfMaxAgeLsa
ospfNbrStateChange
ospfOriginateLsa
ospfTxRetransmit
ospfVirtIfAuthFailure
ospfVirtIfConfigError
ospfVirtIfRxBadPacket
ospfVirtIfStateChange
ospfVirtIfTxRetransmit
ospfVirtNbrStateChange
pethMainPowerUsageOffNotification
pethMainPowerUsageOnNotification
pethPsePortOnOffNotification
pingProbeFailed
pingTestCompleted
pingTestFailed
ptopoConfigChange
risingAlarm
rpMauJabberTrap
sd1cLSStatusChange
sd1cPortStatusChange
topologyChange
traceRoutePathChange
traceRouteTestCompleted
traceRouteTestFailed
vrrpTrapAuthFailure
vrrpTrapNewMaster
warmStart

```

request snmp spoof-trap (Question Mark ?)

```

user@host> request snmp spoof-trap ?
Possible completions:
<trap>                The name of the trap to spoof
ads1AtucInitFailureTrap

```

```
ads1AtucPerfESsThreshTrap
ads1AtucPerfLofsThreshTrap
ads1AtucPerfLolsThreshTrap
ads1AtucPerfLossThreshTrap
ads1AtucPerfLprsThreshTrap
ads1AtucRateChangeTrap
ads1AturPerfESsThreshTrap
ads1AturPerfLofsThreshTrap
ads1AturPerfLossThreshTrap
ads1AturPerfLprsThreshTrap
ads1AturRateChangeTrap
apsEventChannelMismatch
apsEventFEPLF
apsEventModeMismatch
apsEventPSBF
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlsWTrapCircuitDown
dlsWTrapCircuitUp
---(more 10%)---
```

show snmp health-monitor

Syntax	<code>show snmp health-monitor</code> <code><alarms <detail>> <logs></code>
Release Information	Command introduced in Junos OS Release 8.0. Command introduced in Junos OS Release 9.0 for EX Series switches. Statement introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.
Options	none —Display information about all health monitor alarms and logs. alarms <detail> —(Optional) Display detailed information about health monitor alarms. logs —(Optional) Display information about health monitor logs.
Required Privilege Level	view
List of Sample Output	show snmp health-monitor on page 353 show snmp health-monitor alarms detail on page 355
Output Fields	Table 20 on page 351 describes the output fields for the show snmp health-monitor command. Output fields are listed in the approximate order in which they appear.

Table 20: show snmp health-monitor Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
Variable description	Description of the health monitor object instance being monitored.	All levels
Variable name	Name of the health monitor object instance being monitored.	All levels
Value	Current value of the monitored variable in the most recent sample interval.	All levels

Table 20: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <ul style="list-style-type: none"> Alarms: <ul style="list-style-type: none"> active—Entry is fully configured and activated. falling threshold crossed—Value of the variable has crossed the lower threshold limit. rising threshold crossed—Value of the variable has crossed the upper threshold limit. under creation—Entry is being configured and is not yet activated. startup—Alarm is waiting for the first sample of the monitored variable. object not available—Monitored variable of that type is not available to the health monitor agent. instance not available—Monitored variable's instance is not available to the health monitor agent. object type invalid—Monitored variable is not a numeric value. object processing errored—An error occurred when the monitored variable was processed. unknown—State is not one of the above. 	All levels
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of absolute value or delta value .	detail
Startup alarm	<p>Alarm that might be sent when this entry is first activated, depending on the following criteria:</p> <ul style="list-style-type: none"> Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is either falling alarm or rising or falling alarm. Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is rising alarm. Value of the alarm is between the thresholds. 	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.	detail
Creator	Mechanism by which the entry was configured (Health Monitor).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value as a percentage of the maximum possible value.	detail

Table 20: show snmp health-monitor Output Fields (*continued*)

Field Name	Field Description	Level of Output
Falling threshold	Lower limit threshold value as a percentage of the maximum possible value.	detail
Rising event index	Event triggered when the rising threshold is crossed.	detail
Falling event index	Event triggered when the falling threshold is crossed.	detail

Sample Output

show snmp health-monitor

```
user@host> show snmp health-monitor
```

```

Alarm
Index  Variable description                               Value State

32768 Health Monitor: root file system utilization
      jnxHrStoragePercentUsed.1                      58 active

32769 Health Monitor: /config file system utilization
      jnxHrStoragePercentUsed.2                      0 active

32770 Health Monitor: RE 0 CPU utilization
      jnxOperatingCPU.9.1.0.0                        0 active

32773 Health Monitor: RE 0 Memory utilization
      jnxOperatingBuffer.9.1.0.0                     35 active

32775 Health Monitor: jkernel daemon CPU utilization
      Init daemon                                     0 active
      Chassis daemon                                 50 active
      Firewall daemon                                0 active
      Interface daemon                               5 active
      SNMP daemon                                    11 active
      MIB2 daemon                                    42 active
      Sonet APS daemon                               0 active
      VRRP daemon                                    0 active
      Alarm daemon                                   3 active
      PFE daemon                                     0 active
      CRAFT daemon                                   0 active
      Traffic sampling control daemon                 0 active
      Ilmi daemon                                    0 active
      Remote operations daemon                       0 active
      CoS daemon                                     0 active
      Pic Services Logging daemon                     0 active
      Internal Routing Service Daemon                 3 active
      Network Access Service daemon                  0 active
      Forwarding UDP daemon                           0 active
      Routing socket proxy daemon                    0 active
      Disk Monitoring daemon                          1 active
      Inet daemon                                    0 active
      Syslog daemon                                   0 active
      Adaptive Services PIC daemon                   0 active
      ECC parity errors logging Daemon                0 active
      Layer 2 Tunneling Protocol daemon               0 active
      PPPoE daemon                                    3 active

```

	Redundancy device daemon	0 active
	PPP daemon	0 active
	Dynamic Flow Capture Daemon	0 active
32776	Health Monitor: jroute daemon CPU utilization	
	Routing protocol daemon	1 active
	Management daemon	0 active
	Management daemon	0 active
	Command line interface	4 active
	Periodic Packet Management daemon	0 active
	Link Management daemon	0 active
	Pragmatic General Multicast daemon	0 active
	Bidirectional Forwarding Detection daemon	0 active
	SRC daemon	0 active
	audit daemon	0 active
	Event daemon	0 active
32777	Health Monitor: jcrypto daemon CPU utilization	
	IPSec Key Management daemon	0 active
32779	Health Monitor: jkernel daemon Memory utilization	
	Init daemon	47384 active
	Chassis daemon	20204 active
	Firewall daemon	1956 active
	Interface daemon	3340 active
	SNMP daemon	4540 active
	MIB2 daemon	3880 active
	Sonet APS daemon	2632 active
	VRRP daemon	2672 active
	Alarm daemon	1856 active
	PFE daemon	2600 active
	CRAFT daemon	2000 active
	Traffic sampling control daemon	3164 active
	Ilmi daemon	2132 active
	Remote operations daemon	2964 active
	CoS daemon	3044 active
	Pic Services Logging daemon	1944 active
	Internal Routing Service Daemon	1392 active
	Network Access Service daemon	1992 active
	Forwarding UDP daemon	1876 active
	Routing socket proxy daemon	1296 active
	Disk Monitoring daemon	1180 active
	Inet daemon	1296 active
	Syslog daemon	1180 active
	Adaptive Services PIC daemon	3220 active
	ECC parity errors logging Daemon	1100 active
	Layer 2 Tunneling Protocol daemon	3372 active
	PPPoE daemon	1424 active
	Redundancy device daemon	1820 active
	PPP daemon	2060 active
	Dynamic Flow Capture Daemon	10740 active
32780	Health Monitor: jroute daemon Memory utilization	
	Routing protocol daemon	8104 active
	Management daemon	13360 active
	Management daemon	19252 active
	Command line interface	9912 active
	Periodic Packet Management daemon	1484 active
	Link Management daemon	2016 active
	Pragmatic General Multicast daemon	1968 active
	Bidirectional Forwarding Detection daemon	1956 active
	SRC daemon	1772 active

audit daemon	1772 active
Event daemon	1808 active

32781 Health Monitor: jcrypto daemon Memory utilization	
IPSec Key Management daemon	5600 active

show snmp health-monitor alarms detail

```
user@host> show snmp health-monitor alarms detail
```

```
Alarm Index 32768:
  Variable name      jnxHrStoragePercentUsed.1
  Variable OID       1.3.6.1.4.1.2636.3.31.1.1.1.1.1
  Sample type        absolute value
  Startup alarm       rising alarm
  Owner              Health Monitor: root file system
                    utilization
  Creator            Health Monitor
  State              active
  Sample interval     300 seconds
  Rising threshold    80
  Falling threshold   70
  Rising event index  32768
  Falling event index 32768
  Instance Value: 58
  Instance State: active

Alarm Index 32769:
  Variable name      jnxHrStoragePercentUsed.2
  Variable OID       1.3.6.1.4.1.2636.3.31.1.1.1.1.2
  Sample type        absolute value
  Startup alarm       rising alarm
  Owner              Health Monitor: /config file system
                    utilization
  Creator            Health Monitor
  State              active
  Sample interval     300 seconds
  Rising threshold    80
  Falling threshold   70
  Rising event index  32768
  Falling event index 32768
  Instance Value: 0
  Instance State: active

Alarm Index 32770:
  Variable name      jnxOperatingCPU.9.1.0.0
  Variable OID       1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
  Sample type        absolute value
  Startup alarm       rising alarm
  Owner              Health Monitor: RE 0 CPU utilization

  Creator            Health Monitor
  State              active
  Sample interval     300 seconds
  Rising threshold    80
  Falling threshold   70
  Rising event index  32768
  Falling event index 32768
  Instance Value: 0
  Instance State: active
```

Alarm Index 32773:

Variable name jnxOperatingBuffer.9.1.0.0
Variable OID 1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0
Sample type absolute value
Startup alarm rising alarm
Owner Health Monitor: RE 0 Memory utilization

Creator Health Monitor
State active
Sample interval 300 seconds
Rising threshold 80
Falling threshold 70
Rising event index 32768
Falling event index 32768
Instance Value: 35
Instance State: active

Alarm Index 32775:

Variable name sysAppElmtRunCPU.3
Variable OID 1.3.6.1.2.1.54.1.2.3.1.9.3
Sample type delta value
Startup alarm rising alarm
Owner Health Monitor: jkernel daemon CPU utilization

Creator Health Monitor
State active
Sample interval 300 seconds
Rising threshold 24000
Falling threshold 21000
Rising event index 32768
Falling event index 32768
Instance Name: sysAppElmtRunCPU.3.1.1
Instance Description: Init daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElmtRunCPU.3.2.2786
Instance Description: Chassis daemon
Instance Value: 50
Instance State: active

Instance Name: sysAppElmtRunCPU.3.3.2938
Instance Description: Firewall daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElmtRunCPU.3.4.2942
Instance Description: Interface daemon
Instance Value: 5
Instance State: active

Instance Name: sysAppElmtRunCPU.3.7.7332
Instance Description: SNMP daemon
Instance Value: 11
Instance State: active

Instance Name: sysAppElmtRunCPU.3.9.2914
Instance Description: MIB2 daemon
Instance Value: 42

```
Instance State: active

Instance Name: sysAppElmtRunCPU.3.12.2916
Instance Description: Sonet APS daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElmtRunCPU.3.13.2917
Instance Description: VRRP daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElmtRunCPU.3.14.2787
Instance Description: Alarm daemon
Instance Value: 3
Instance State: active

Instance Name: sysAppElmtRunCPU.3.15.2940
Instance Description: PFE daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElmtRunCPU.3.16.2788
Instance Description: CRAFT daemon
Instance Value: 0
Instance State: active

Instance Name: sysAppElmtRunCPU.3.17.2918
Instance Description: Traffic sampling control daemon
---(more 23%)---
```

show snmp inform-statistics

Syntax	show snmp inform-statistics
Release Information	Command introduced in Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about Simple Network Management Protocol (SNMP) inform requests.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show snmp inform-statistics on page 358
Output Fields	Table 21 on page 358 describes the output fields for the show snmp inform-statistics command. Output fields are listed in the approximate order in which they appear.

Table 21: show snmp inform-statistics Output Fields

Field Name	Field Description
Target Name	Name of the device configured to receive and respond to SNMP informs.
Address	IP address of the target device.
Sent	Number of informs sent to the target device and acknowledged by the target device.
Pending	Number of informs held in memory pending a response from the target device.
Discarded	Number of informs discarded after the specified number of retransmissions to the target device were attempted.
Timeouts	Number of informs that did not receive an acknowledgement from the target device within the timeout specified.
Probe Failures	Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address).

Sample Output

show snmp inform-statistics

```

user@host> show snmp inform-statistics
Inform Request Statistics:
  Target Name: TA1_v3_md5_none Address: 172.17.20.184
    Sent: 176, Pending: 0
    Discarded: 0, Timeouts: 0, Probe Failures: 0
  Target Name: TA2_v3_sha_none Address: 192.168.110.59

```

Sent: 0, Pending: 4
Discarded: 84, Timeouts: 0, Probe Failures: 258
Target Name: TA5_v2_none Address: 172.17.20.184
Sent: 0, Pending: 0
Discarded: 2, Timeouts: 10, Probe Failures: 0

show snmp mib

Syntax	<code>show snmp mib (get get-next walk) (ascii decimal) <i>object-id</i></code>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>ascii and decimal options introduced in Junos OS Release 9.6.</p> <p>ascii and decimal options introduced in Junos OS Release 9.6 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.
Options	<p>get—Retrieve and display one or more SNMP object values.</p> <p>get-next—Retrieve and display the next SNMP object values.</p> <p>walk—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.</p> <p>ascii—Display the SNMP object's string indices as an ASCII-key representation.</p> <p>decimal—Display the SNMP object values in the decimal (default) format. The decimal option is the default option for this command. Therefore, issuing the show snmp mib (get get-next walk) decimal object-id and the show snmp mib (get get-next walk) object-id commands display the same output.</p> <p>object-id—The object can be represented by a sequence of dotted integers (such as 1.3.6.1.2.1.2) or by its subtree name (such as interfaces). When entering multiple objects, enclose the objects in quotation marks.</p>
Required Privilege Level	snmp—To view this statement in the configuration.
List of Sample Output	<p>show snmp mib get on page 361</p> <p>show snmp mib get (Multiple Objects) on page 361</p> <p>show snmp mib get (Layer 2 Policer) on page 361</p> <p>show snmp mib get-next on page 361</p> <p>show snmp mib get-next (Specify an OID) on page 361</p> <p>show snmp mib walk on page 361</p> <p>show snmp mib walk (QFX Series) on page 361</p> <p>show snmp mib walk decimal on page 362</p> <p>show snmp mib walk (ASCII) on page 362</p> <p>show snmp mib walk (Multiple Indices) on page 362</p> <p>show snmp mib walk decimal (Multiple Indices) on page 362</p>
Output Fields	Table 22 on page 361 describes the output fields for the show snmp mib command. Output fields are listed in the approximate order in which they appear.

Table 22: show snmp mib Output Fields

Field Name	Field Description
<i>name</i>	Object name and numeric instance value.
<i>object value</i>	Object value. The Junos OS translates OIDs into the corresponding object names.

Sample Output

show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get ?sysObjectID.0 sysUpTime.0?
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

show snmp mib get (Layer 2 Policer)

```
user@host> show snmp mib get ifInOctets.25970
ifInOctets.25970 = 7545720
```

show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
Networks, Inc.
```

show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. m20 internet router, kernel
Junos OS Release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004
Juniper Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

show snmp mib walk (QFX Series)

```
user@switch> show snmp mib walk system
sysDescr.0 = Juniper Networks, Inc. qfx3500s internet router, kernel JUNOS
11.1-20100926.0 #0: 2010-09-26 06:17:38 UTC Build date: 2010-09-26 06:00:10
sysObjectID.0 = jnxProductQFX3500
sysUpTime.0 = 138980301
sysContact.0 = System Contact
```

```
sysName.0      = LabQFX3500
sysLocation.0  = Lab
sysServices.0  = 4
```

show snmp mib walk decimal

```
user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100
```

show snmp mib walk (ASCII)

```
show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100
```

show snmp mib walk (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```

show snmp mib walk decimal (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.....
```


show snmp rmon

Syntax	show snmp rmon <alarms <brief detail> events <brief detail> logs>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms and events.
Options	<p>none—Display information about all RMON alarms and events.</p> <p>alarms—(Optional) Display information about RMON alarms.</p> <p>brief detail—(Optional) Display brief or detailed information about RMON alarms or events.</p> <p>events—(Optional) Display information about RMON events.</p> <p>logs—(Optional) Display information about RMON monitoring logs.</p>
Required Privilege Level	view
List of Sample Output	show snmp rmon on page 365 show snmp rmon alarms detail on page 365 show snmp rmon events detail on page 366
Output Fields	Table 23 on page 363 describes the output fields for the show snmp rmon command. Output fields are listed in the approximate order in which they appear.

Table 23: show snmp rmon Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels

Table 23: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
State	<p>State of the alarm or event entry:</p> <p>Alarms:</p> <ul style="list-style-type: none"> • active—Entry is fully configured and activated. • falling threshold crossed—Value of the variable has crossed the lower threshold limit. • rising threshold crossed—Value of the variable has crossed the upper threshold limit. • under creation—Entry is being configured and is not yet activated. • startup—Alarm is waiting for the first sample of the monitored variable. • object not available—Monitored variable of that type is not available to the SNMP agent. • instance not available—Monitored variable's instance is not available to the SNMP agent. • object type invalid—Monitored variable is not a numeric value. • object processing errored—An error occurred when the monitored variable was processed. • unknown—State is not one of the above. <p>Events:</p> <ul style="list-style-type: none"> • active—Entry has been fully configured and activated. • under creation—Entry is being configured and is not yet activated. • unknown—State is not one of the above. 	All levels
Variable name	Name of the SNMP object instance being monitored.	All levels
Event Index	Event identifier.	All levels
Type	<p>Type of notification made when an event is triggered. It can be one of the following:</p> <ul style="list-style-type: none"> • log—A system log message is generated and an entry is made to the log table. • snmptrap—An SNMP trap is sent to the configured destination. • log and trap—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination. • none—Neither log nor trap will be sent. 	detail
Last Event	Date and time of the last event. It has the format <i>yyyy-mm-dd hh:mm:ss timezone</i> .	brief
Community	Identifies the trap group used for sending the SNMP trap.	detail
Variable OID	Object ID to which the variable name is resolved. The format is x.x.x.x.	detail
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of absolute value or delta value .	detail

Table 23: show snmp rmon Output Fields (*continued*)

Field Name	Field Description	Level of Output
Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria: <ul style="list-style-type: none"> Alarm is sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is either rising alarm or rising or falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is either falling alarm or rising or falling alarm. Alarm is <i>not</i> sent when one of the following situations exists: <ul style="list-style-type: none"> Value of the alarm is above or equal to the rising threshold and the startup type is falling alarm. Value of the alarm is below or equal to the falling threshold and the startup type is rising alarm. Value of the alarm is between the thresholds. 	detail
Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has monitor prepended to it.	detail
Creator	Mechanism by which the entry was configured (CLI or SNMP).	detail
Sample interval	Time period between samples (in seconds).	detail
Rising threshold	Upper limit threshold value configured by the user.	detail
Falling threshold	Lower limit threshold value configured by the user.	detail
Rising event index	Event triggered when the rising threshold is crossed.	detail
Falling event index	Event triggered when the falling threshold is crossed.	detail
Current value	Current value of the monitored variable in the most recent sample interval.	detail

Sample Output

show snmp rmon

```

user@host> show snmp rmon
Alarm
Index  State                               Variable name
  1    falling threshold crossed         ifInOctets.1

Event
Index  Type                                Last Event
  1    log and trap                     2002-01-30 01:13:01 PST

```

show snmp rmon alarms detail

```

user@host> show snmp rmon alarms detail

```

```
Alarm Index 1:
Variable name      ifInOctets.1
Variable OID       1.3.6.1.2.1.2.2.1.10.1
Sample type        delta value
Startup alarm      rising or falling alarm
Owner              monitor
Creator            CLI
State              falling threshold crossed
Sample interval    60 seconds
Rising threshold   100000
Falling threshold  80000
Rising event index 1
Falling event index 1
Current value      0
```

show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
Type          log and trap
Community     boy-elroy
Last event    2002-01-30 01:13:01 PST
Creator       CLI
State         active
```

show snmp rmon history

Syntax	show snmp rmon history <i><history-index></i> <i><sample-index></i>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the contents of the RMON history group.
Options	<p>none—Display all the entries in the RMON history group.</p> <p>history-index—(Optional) Display the contents of the specified entry in the RMON history group.</p> <p>sample-index—(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> clear snmp rmon history on page 342
List of Sample Output	show snmp rmon history 1 on page 368 show snmp rmon history 1 sample 15 on page 369
Output Fields	<p>Table 24 on page 367 lists the output fields for the show smp rmon history command. Output fields are listed in the approximate order in which they appear.</p>

Table 24: show smp rmon history Output Fields

Field Name	Field Description
History Index	Identifies this RMON history entry within the RMON history group.
Owner	The entity that configured this entry. Range is 0 to 32 alphanumeric characters.
Status	The status of the RMON history entry.
Interface or Data Source	The ifindex object that identifies the interface that is being monitored.
Interval	The interval (in seconds) configured for this RMON history entry.
Buckets Requested	The requested number of buckets (intervals) configured for this RMON history entry.
Buckets Granted	The number of buckets granted for this RMON history entry.

Table 24: show smp rmon history Output Fields (*continued*)

Field Name	Field Description
Sample Index	<p>The sample statistics taken at the specified interval.</p> <ul style="list-style-type: none"> • Drop Events—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. • Octets—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • Packets—Total number of packets. • Broadcast Packets—Number of broadcast packets. • Multicast Packets—Number of multicast packets. • CRC 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). • Undersize Pkts—Number of packets received during this sampling interval that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed. • Oversize Pkts—Number of packets received during the sampling interval that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed. • Fragments—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. • Jabbers—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. • 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. • Utilization(%)—The best estimate of the mean physical layer network utilization on this interface during this sampling interval, in hundredths of a percent.

Sample Output

show snmp rmon history 1

```

user@host> show snmp rmon history 1
History Index 1:
Interface                171
Requested Buckets        50
Interval                 10

Sample Index 1: Interval Start: Tue Feb 12 04:12:32 2008
Drop Events              0
Octets                   486
Packets                  2

```

Broadcast Packet	0
Multicast Packets	2
CRC errors	0
Undersize Pkts	0
Oversize Pkts	0
Fragments	0
Jabbers	0
Collisions	0
Utilization(%)	0

Sample Index 2: Interval Start: Tue Feb 12 04:12:42 2008

Drop Events	0
Octets	486
Packets	2
Broadcast Packet	0
Multicast Packets	2
CRC errors	0
Undersize Pkts	0
Oversize Pkts	0
Fragments	0
Jabbers	0
Collisions	0
Utilization(%)	0

Sample Index 3: Interval Start: Tue Feb 12 04:12:52 2008

Drop Events	0
Octets	486
Packets	2
Broadcast Packet	0
Multicast Packets	2
CRC errors	0
Undersize Pkts	0
Oversize Pkts	0
Fragments	0
Jabbers	0
Collisions	0
Utilization(%)	0

show snmp rmon history 1 sample 15

user@host> show snmp rmon history 1 sample 15

Index 1
 Owner = monitor
 Status = valid
 Data Source = ifIndex.17
 Interval = 1800
 Buckets Requested = 50
 Buckets Granted = 50

Sample Index 44: Interval Start: Thu Jan 1 00:08:35 1970

Drop Events	= 0
Octets	= 0
Packets	= 0
Broadcast Pkts	= 0
Multicast Pkts	= 0
CRC Errors	= 0
Undersize Pkts	= 0
Oversize Pkts	= 0
Fragments	= 0
Jabbers	= 0

```
Collisions = 0  
Utilization (%) = 0
```


show snmp statistics

Syntax	show snmp statistics
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.
Options	This command has no options.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear snmp statistics on page 343
List of Sample Output	show snmp statistics on page 374
Output Fields	Table 25 on page 371 describes the output fields for the show snmp statistics command. Output fields are listed in the approximate order in which they appear.

Table 25: show snmp statistics Output Fields

Field Name	Field Description
Input	<p>Information about received packets:</p> <ul style="list-style-type: none"> • Packets(snmplnPkts)—Total number of messages delivered to the SNMP entity from the transport service. • Bad versions—(snmplnBadVersions) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • Bad community names—(snmplnBadCommunityNames) Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • Bad community uses—(snmplnBadCommunityUses) Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • ASN parse errors—(snmplnASNParseErrs) Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • Too big—(snmplnTooBig) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of tooBig. • No such names—(snmplnNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmplnBadValues) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of badValue. • Read only—(snmplnReadOnly) Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of readOnly. Only incorrect implementations of SNMP generate this error.

Table 25: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Input (continued)	<ul style="list-style-type: none"> • General errors—(snmpInGenErrs) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of genErr. • Total requests varbinds—(snmpInTotalReqVars) Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP GetRequest and GetNext PDUs. • Total set varbinds—(snmpInSetVars) Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP SetRequest PDUs. • Get requests—(snmpInGetRequests) Total number of SNMP GetRequest PDUs that have been accepted and processed by the SNMP entity. • Get nexts—(snmpInGetNexts) Total number of SNMP GetNext PDUs that have been accepted and processed by the SNMP entity. • Set requests—(snmpInSetRequests) Total number of SNMP SetRequest PDUs that have been accepted and processed by the SNMP entity. • Get responses—(snmpInGetResponses) Total number of SNMP GetResponse PDUs that have been accepted and processed by the SNMP entity. • Traps—(snmpInTraps) Total number of SNMP traps generated by the SNMP entity. • Silent drops—(snmpSilentDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • Proxy drops—(snmpProxyDrops) Total number of GetRequest, GetNextRequest, GetBulkRequest, SetRequests, and InformRequest PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned. • Commit pending drops—Number of SNMP packets for Set requests dropped because of a previous pending SNMP Set request on the committed configuration. • Throttle drops—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 25: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
V3 Input	<p>Information about SNMP version 3 packets:</p> <ul style="list-style-type: none"> • Unknown security models—(snmpUnknownSecurityModels) Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • Invalid messages—(snmpInvalidMsgs) Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • Unknown pdu handlers—(snmpUnknownPDUHandlers) Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • Unavailable contexts—(snmpUnavailableContexts) Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • Unknown contexts—(snmpUnknownContexts) Total number of requests received for a context that is unknown to the SNMP engine. • Unsupported security levels—(usmStatsUnsupportedSecLevels) Total number of packets received by the SNMP engine that were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • Not in time windows—(usmStatsNotInTimeWindows) Total number of packets received by the SNMP engine that were dropped because they appeared outside the authoritative SNMP engine's window. • Unknown user names—(usmStatsUnknownUserNames) Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • Unknown engine ids—(usmStatsUnknownEngineIDs) Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine. • Wrong digests—(usmStatsWrongDigests) Total number of packets received by the SNMP engine that were dropped because they did not contain the expected digest value. • Decryption errors—(usmStatsDecryptionErrors) Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.

Table 25: show snmp statistics Output Fields (*continued*)

Field Name	Field Description
Output	<p>Information about transmitted packets:</p> <ul style="list-style-type: none"> • Packets—(snmpOutPkts) Total number of messages passed from the SNMP entity to the transport service. • Too big—(snmpOutTooBig) Total number of SNMP PDUs generated by the SNMP entity with an error status field of tooBig. • No such names—(snmpOutNoSuchNames) Total number of SNMP PDUs delivered to the SNMP entity with an error status field of noSuchName. • Bad values—(snmpOutBadValues) Total number of SNMP PDUs generated by the SNMP entity with an error status field of badValue. • General errors—(snmpOutGenErrs) Total number of SNMP PDUs generated by the SNMP entity with an error status field of genErr. • Get requests—(snmpOutGetRequests) Total number of SNMP GetRequest PDUs generated by the SNMP entity. • Get nexts—(snmpOutGetNexts) Total number of SNMP GetNext PDUs generated by the SNMP entity. • Set requests—(snmpOutSetRequests) Total number of SNMP SetRequest PDUs generated by the SNMP entity. • Get responses—(snmpOutGetResponses) Total number of SNMP GetResponse PDUs generated by the SNMP entity. • Traps—(snmpOutTraps) Total number of SNMP traps generated by the SNMP entity.

Sample Output

show snmp statistics

```

user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too big: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too big: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0

```

show snmp v3

Syntax	<code>show snmp v3</code> <code><access <brief detail> community general groups notify <filter> target <address parameters> users></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.
Options	<p>none—Display all of the SNMPv3 operating configuration.</p> <p>access—(Optional) Display SNMPv3 access information.</p> <p>brief detail—(Optional) Display brief or detailed information about SNMPv3 access information.</p> <p>community—(Optional) Display SNMPv3 community information.</p> <p>general—(Optional) Display SNMPv3 general information.</p> <p>groups—(Optional) Display SNMPv3 security-to-group information.</p> <p>notify <filter>—(Optional) Display SNMPv3 notify and, optionally, notify filter information.</p> <p>target <address parameters>—(Optional) Display SNMPv3 target and, optionally, either target address or target parameter information.</p> <p>users—(Optional) Display SNMPv3 user information.</p>
Additional Information	To edit the default display of the show snmp v3 command, specify options in the show statement at the [edit snmp v3] hierarchy level.
Required Privilege Level	view
List of Sample Output	show snmp v3 on page 377
Output Fields	Table 26 on page 376 describes the output fields for the show snmp v3 command. Output fields are listed in the approximate order in which they appear.

Table 26: show snmp v3 Output Fields

Field Name	Field Description
Access control	<p>Information about access control:</p> <ul style="list-style-type: none"> • Group—Group name for which the configured access privileges apply. The group, together with the context prefix and the security model and security level, forms the index for this table. • Context prefix—SNMPv3 context for which the configured access privileges apply. • Security model/level—Security model and security level for which the configuration access privileges apply. • Read view—Identifies the MIB view applied to SNMPv3 read operations. • Write view—Identifies the MIB view applied to SNMPv3 write operations. • Notify view—Identifies the MIB view applied to outbound SNMP notifications.
Engine	<p>Information about local engine configuration:</p> <ul style="list-style-type: none"> • Local engine ID—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • Engine boots—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • Engine time—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • Max msg size—Maximum message size the sender can accommodate.
Engine ID	<p>Information about engine ID:</p> <ul style="list-style-type: none"> • Local engine ID—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • Engine boots—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • Engine time—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • Max msg size—Maximum message size the sender can accommodate. • Engine ID—SNMPv3 engine ID associated with each user. • User—SNMPv3 user. • Auth/Priv—Authentication and encryption algorithm available for use by each user. • Storage—Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status. • Status—Status of the conceptual row. Only rows with an active status are used by the SNMPv3 engine.
Group name	Name of the group to which this entry belongs.
Security model	Identifies the security model context for the security name.
Security name	Used with the security model; identifies a specific security name instance. Each security model/security name combination can be assigned to a specific group.
Storage type	Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.
Status	Status of the conceptual row. Only rows with active status are used by the SNMPv3 engine.

Sample Output

show snmp v3

```

user@host> show snmp v3
Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots:      38
Engine time:       64583 seconds
Max msg size:      2048 bytes

Engine ID: local
  User          Auth/Priv  Storage  Status
  user1         md5/des   nonvolatile active
  user2         sha/none  nonvolatile active
  user3         none/none nonvolatile active

Engine ID: 81 00 0a 4c 04 64 64 64 64
  User          Auth/Priv  Storage  Status
  UNEW         md5/none  nonvolatile active
Group name      Security model  Security name      Storage type      Status
g1             usm          user1              nonvolatile active
g2             usm          user2              nonvolatile active
g3             usm          user3              nonvolatile active

Access control:
Group           Context prefix  Security model/level  Read view  Write view  Notify view
g1              usm/privacy  v1                   v1
g2              usm/authent  v1                   v1
g3              usm/none     v1                   v1

```


CHAPTER 19

Operational Commands: Port Mirroring

- `show analyzer`

show analyzer

Syntax	show analyzer <i>analyzer-name</i>
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches.
Description	Display information about analyzers configured for mirroring.
Options	<i>analyzer-name</i> —(Optional) Displays the status of a specific analyzer on the switch.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding Port Mirroring on EX Series Switches on page 9
List of Sample Output	show analyzer on page 380
Output Fields	Table 27 on page 380 lists the output fields for the command-name command. Output fields are listed in the approximate order in which they appear.

Table 27: show analyzer Output Fields

Field Name	Field Description
Analyzer name	Displays the name of the analyzer.
Output interface	Specifies a local interface to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.
Output VLAN	Specifies a VLAN to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.
Mirror ratio	Displays the ratio of packets to be mirrored.
Egress monitored interfaces	Displays interfaces for which traffic exiting the interfaces is mirrored.
Ingress monitored interfaces	Displays interfaces for which traffic entering the interfaces is mirrored.
Ingress monitored VLANs	Displays VLANs for which traffic entering the VLAN is mirrored.

Sample Output

show analyzer

```

user@host> show analyzer
Analyzer name       : employee-monitor
Output interface    : ge-0/0/10.0
Output VLAN         : remote-analyzer
Mirror ratio        : 1
Loss priority       : High
Egress monitored interfaces : ge-0/0/3.0

```

```
Ingress monitored interfaces : ge-0/0/0.0  
Ingress monitored interfaces : ge-0/0/1.0
```


CHAPTER 20

Operational Commands: sFlow

- `show sflow`
- `show sflow interface`
- `show sflow collector`
- `clear sflow collectors statistics`

show sflow

Syntax	show sflow <collector> <interface>
Release Information	Command introduced in Junos OS Release 9.3 for EX Series switches.
Description	Display default sFlow technology configuration information.
Options	<p>none—Display default sFlow technology configuration information.</p> <p>collector—(Optional) Display standard status information about the specified sFlow collector.</p> <p>interface—(Optional) Display standard status information about the specified sFlow interface.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show sflow interface on page 386 • show sflow collector on page 388 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84
List of Sample Output	show sflow on page 385
Output Fields	Table 28 on page 384 lists the output fields for the show sflow command. Output fields are listed in the approximate order in which they appear.

Table 28: show sflow Output Fields

Field Name	Field Description	Level of Output
sFlow	Status of the feature: enabled or disabled .	All levels
Sample rate egress	Rate at which egress packets are sampled.	All levels
Sample rate ingress	Rate at which ingress packets are sampled.	All levels
Sample limit	Number of packets sampled per second. The sampling limit cannot be configured and is set to 300 packets per second.	All levels
Polling interval	Interval at which the sFlow agent polls the interface.	All levels
Agent ID	The IP address assigned to the sFlow agent.	All levels

Table 28: show sflow Output Fields (*continued*)

Field Name	Field Description	Level of Output
Source IP address	The IP address for the sFlow datagram.	All levels

Sample Output

show sflow

```
sFlow          : Enabled
Sample rate egress  : 1:1000
Sample rate ingress : 1: 2048: Disabled
Sample limit       : 300 packets/second
Polling interval   : 20 seconds
Agent ID           : 10.93.54.7
Source IP address  : 10.93.54.7
```

show sflow interface

Syntax	show sflow interface
Release Information	Command introduced in Junos OS Release 9.3 for EX Series switches.
Description	Display the interfaces on which sFlow technology is enabled and the sampling parameters.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show sflow on page 384 • show sflow collector on page 388 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84
List of Sample Output	show sflow interface on page 386
Output Fields	Table 29 on page 386 lists the output fields for the show sflow interface command. Output fields are listed in the approximate order in which they appear.

Table 29: show sflow interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interfaces on which sFlow technology is enabled.	All levels
Status Egress	Indicates whether egress sampling rate is enabled.	All levels
Status Ingress	Indicates whether ingress sampling rate is enabled.	All levels
Sample rate Egress	Rate at which egress packets are sampled.	All levels
Sample rate Ingress	Rate at which ingress packets are sampled.	All levels
Adapted sample rate Egress	Adapted rate at which egress packets are sampled.	All levels
Adapted sample rate Ingress	Adapted rate at which ingress packets are sampled.	All levels
Polling-interval	The interval at which the sFlow agent polls the interface.	All levels

Sample Output

show sflow interface

```
Interface          Status      Sample rate  Adapted sample rate  Polling-interval
```


ge-0/0/0.0	Egress Enabled	Ingress Disabled	Egress 1000	Ingress 2048	Egress 1000	Ingress 2048	20
------------	-------------------	---------------------	----------------	-----------------	----------------	-----------------	----

show sflow collector

Syntax	show sflow collector
Release Information	Command introduced in Junos OS Release 9.3 for EX Series switches.
Description	Display a list of configured sFlow collectors and their properties.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show sflow on page 384 • show sflow interface on page 386 • Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 33 • Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84
List of Sample Output	show sflow collector on page 388
Output Fields	Table 30 on page 388 lists the output fields for the show sflow collector command. Output fields are listed in the approximate order in which they appear.

Table 30: show sflow collector Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the collector.	All levels
UDP port	UDP port number.	All levels
No of samples	Number of samples sent to collector.	All levels

Sample Output

show sflow collector

```

IP-address    UDP-Port    No of samples
10.204.32.46  5600        1000
100.204.32.76 3400        1000

```

clear sflow collectors statistics

Syntax	clear sflow collectors statistics
Release Information	Command introduced in JUNOS Release 9.5 for EX Series switches.
Description	Clear the sFlow collector's statistics.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show sflow collector on page 388• Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 84

Sample Output

```
clear sflow collectors statistics
```


CHAPTER 21

Operational Commands: Ethernet OAM Connectivity Fault Management

- `clear oam ethernet connectivity-fault-management delay-statistics`
- `clear oam ethernet connectivity-fault-management sla-iterator-statistics`
- `clear oam ethernet connectivity-fault-management statistics`
- `monitor ethernet delay-measurement`
- `show oam ethernet connectivity-fault-management delay-statistics`
- `show oam ethernet connectivity-fault-management forwarding-state`
- `show oam ethernet connectivity-fault-management interfaces`
- `show oam ethernet connectivity-fault-management path-database`
- `show oam ethernet connectivity-fault-management mep-database`
- `show oam ethernet connectivity-fault-management mip`
- `show oam ethernet connectivity-fault-management sla-iterator-statistics`

clear oam ethernet connectivity-fault-management delay-statistics

Syntax	<code>clear oam ethernet connectivity-fault-management delay-statistics maintenance-association <i>maintenance-association-name</i> maintenance-domain <i>maintenance-domain-name</i> <logical-system <i>logical-system-name</i>> <one-way> <two-way></code>
Release Information	Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 11.4 for EX Series switches.
Description	On MX Series routers and EX Series switches, clear ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM) delay statistics and ETH-DM frame counts.
Options	<p><code>maintenance-association <i>maintenance-association-name</i></code>—Clear ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association.</p> <p><code>maintenance-domain <i>maintenance-domain-name</i></code>—Clear ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance domain.</p> <p><code>logical-system <i>logical-system-name</i></code>—(MX Series routers only) (Optional) Clear ETH-DM delay statistics and ETH-DM frame counts for the specified logical system.</p> <p><code>one-way</code>—(Optional) Clear one-way ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association, maintenance domain, or (on the routers only) logical system.</p> <p><code>two-way</code>—(Optional) Clear two-way ETH-DM delay statistics and ETH-DM frame counts for the specified maintenance association, maintenance domain, or (on the routers only) logical system.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• clear oam ethernet connectivity-fault-management statistics• show oam ethernet connectivity-fault-management delay-statistics on page 401• show oam ethernet connectivity-fault-management interfaces
List of Sample Output	clear oam ethernet connectivity-fault-management delay statistics on page 392
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management delay statistics

```
user@switch> clear oam ethernet connectivity-fault-management delay-statistics  
maintenance-domain md1 maintenance-association ma1  
Delay statistics entries cleared
```


clear oam ethernet connectivity-fault-management sla-iterator-statistics

Syntax	<code>clear oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-association <i>maintenance-association-name</i> maintenance-domain <i>maintenance-domain-name</i> <local-mep <i>local-mep-id</i>> <remote-mep <i>remote-mep-id</i>> sla-iterator <i>sla-iterator</i></code>
Release Information	Command introduced in Junos OS Release 11.4 for EX Series switches. Command introduced in Junos OS Release 13.2 for MX Series routers.
Description	Clear Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics. For MX Series routers, clear the SLA iterator statistics and proactive Ethernet synthetic loss measurement (ETH-SLM) statistics.
Options	<code>maintenance-association <i>maintenance-association-name</i></code> —Name of the maintenance association. <code>maintenance-domain <i>maintenance-domain-name</i></code> —Name of the maintenance domain. <code>local-mep <i>local-mep-id</i></code> —(Optional) Identifier of the local MEP. <code>remote-mep <i>remote-mep-id</i></code> —(Optional) Identifier of the remote MEP. <code>sla-iterator <i>sla-iterator</i></code> — Name of the SLA iterator profile.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none">• Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102
List of Sample Output	clear oam ethernet connectivity-fault-management sla-iterator- statistics on page 394
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

`clear oam ethernet connectivity-fault-
-management sla-iterator- statistics`

```
user@switch> clear oam ethernet connectivity-fault-management sla-iterator-statistics  
maintenance-domain md1 maintenance-association ma1 local-mep 1 remote-mep 2 sla-iterator  
i1  
Iterator statistics entries cleared
```


clear oam ethernet connectivity-fault-management statistics

Syntax	<code>clear oam ethernet connectivity-fault-management statistics</code> <code><interface ethernet-interface-name></code> <code><level md-level></code>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Clear all statistics maintained by CFM.
Options	<code>interface ethernet-interface-name</code> —(Optional) Clear CFM statistics only for MEPs attached to the specified Ethernet physical interface. <code>level level</code> —(Optional) Clear CFM statistics only for MEPs within CFM maintenance domains (MDs) of the specified level.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show oam ethernet connectivity-fault-management interfaces on page 409• show oam ethernet connectivity-fault-management path-database on page 415• show oam ethernet connectivity-fault-management mip on page 424
List of Sample Output	clear oam ethernet connectivity-fault-management statistics on page 395
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear oam ethernet connectivity-fault-management statistics

```
user@host> clear oam ethernet connectivity-fault-management statistics
Cleared statistics of all CFM sessions
```

monitor ethernet delay-measurement

Syntax `monitor ethernet delay-measurement maintenance-domain md-name
maintenance-association ma-name (one-way | two-way) (remote-mac-address | mep
remote-mep-id) <count count> <no-session-id-tlv> <priority 802.1p value> <size size>
<wait time>`

Release Information Command introduced in Junos OS Release 11.4 for EX Series switches.

Description Start an ITU-T Y.1731 Ethernet frame delay measurement session between the specified local connectivity fault management (CFM) maintenance association end point (MEP) and the specified remote MEP, and display a summary of the frames exchanged in the measurement session. Frame delay measurement statistics are stored at one of the MEPs for later retrieval.



NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type Ctrl +C to explicitly quit the `monitor ethernet delay-measurement` command and return to the CLI command prompt.

To start an Ethernet frame delay measurement session, the switch initiates an exchange of frames carrying one-way or two-way frame delay measurement protocol data units (PDUs) between the local and remote MEPs. The frame counts—the types of and number of Ethernet frame delay measurement PDU frames exchanged to measure frame delay times—are displayed as the run-time output of the `monitor ethernet delay-measurement` command and are also stored at both the initiator and receiver MEPs for later retrieval. Ethernet frame delay measurement statistics, described below, are measured and stored at only one of the MEPs:

Frame delay—The difference, in microseconds, between the time a frame is sent and when it is received.

Frame delay variation—The difference, in microseconds, between consecutive frame delay values. Frame delay variation is sometimes called “frame jitter.”

For one-way Ethernet frame delay measurement, only the receiver MEP (on the remote system) collects statistics. For two-way Ethernet frame delay measurement, only the initiator MEP (on the local system) collects statistics.

Options `count count`—(Optional) Number of frames to send to the specified peer MEP. The range of values is 1 through 65,535 frames. The default value is 10 frames.

`maintenance-association ma-name`—Name of an existing CFM maintenance association.

`maintenance-domain md-name`—Name of an existing CFM maintenance domain.

`mep remote-mep-id`—Numeric identifier of the peer MEP with which to perform Ethernet frame delay measurement. The discovered MAC address of the peer MEP is used. The range of values is 1 through 8191.

no-session-id-tlv—(Optional) Prevent insertion of the session ID TLV in the request frame.

one-way—Measurement type is one-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a one-way delay measurement request (IDM) frame and the time at which the receiver MEP receives the frame.

priority 802.1p value—(Optional) Priority of the delay measurement request frame supported by both one-way delay measurement and two-way delay measurement. The range of values is from 0 through 7. The default value is zero.

remote-mac-address—Unicast MAC address of the peer MEP with which to perform Ethernet frame delay measurement. Specify the MAC address as six hexadecimal bytes in *nn:nn:nn:nn:nn:nn* format. Multicast MAC addresses are not supported.

size size —(Optional) Size of the data TLV to be included in the request frame. The range of values is from 1 through 1400 bytes.

two-way—Measurement type is two-way Ethernet frame delay measurement, which is based on the difference between the time at which the initiator MEP sends a two-way delay measurement message (DMM) frame and the time at which the initiator MEP receives an associated two-way delay measurement reply (DMR) frame from the responder MEP, subtracting the time elapsed at the responder MEP.

wait time—(Optional) Number of seconds to wait between sending frames. The range of values is from 1 through 255 seconds. The default value is 1 second.

Required Privilege Level trace and maintenance

Related Documentation

- [Configuring an Iterator Profile on a Switch \(CLI Procedure\) on page 102](#)
- [show oam ethernet connectivity-fault-management mep-database on page 418](#)
- [show oam ethernet connectivity-fault-management mep-statistics](#)
- [show oam ethernet connectivity-fault-management delay-statistics on page 401](#)
- [clear oam ethernet connectivity-fault-management statistics on page 395](#)

List of Sample Output

- [monitor ethernet delay-measurement one-way on page 399](#)
- [monitor ethernet delay-measurement two-way on page 399](#)
- [monitor ethernet delay-measurement two-way \(Invalid DMR Frames Received\) on page 399](#)

Output Fields The **monitor ethernet delay-measurement** command displays different output at the CLI, depending on whether you start a one-way or two-way frame delay measurement:

- [Table 31 on page 398](#) lists the run-time output fields for the **monitor ethernet delay-measurement one-way** command.
- [Table 32 on page 398](#) lists the run-time output fields for the **monitor ethernet delay-measurement two-way** command.

Output fields are listed in the approximate order in which they appear.

Table 31: monitor ethernet delay-measurement one-way Output Fields

Output Field Name	Output Field Description
One-way ETH-DM request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
IDM Frames sent	PDU frames sent to the remote MEP in this ETH-DM session.
Packets transmitted	Total number of IDM PDU frames sent to the remote MEP during this measurement session.
Average delay	Average two-way frame delay measured in this session.
Average delay variation	Average frame jitter measured in this session.
Best case delay	Lowest two-way frame delay measured in this session.
Worst case delay	Highest two-way frame delay measured in this session.

NOTE: For one-way delay measurement, these CLI output fields display **NA** ("not applicable") at the initiator MEP because one-way frame delay measurements occur at the receiver MEP.

Table 32: monitor ethernet delay-measurement two-way Output Fields

Output Field Name	Output Field Description
Two-way Ethernet frame delay measurement request to	Unicast MAC address of the remote peer MEP.
Interface	Name of the Ethernet physical, logical, or trunk interface to which the local MEP is attached.
DMR received from	Unicast MAC address of the remote MEP that transmitted this DMR frame in response to a DMM frame.
Delay	Two-way delay, in microseconds, for the initiator-transmitted DMM frame.
Delay variation	Difference, in microseconds, between the current and previous delay values. This is also known as <i>jitter</i> .
Packets transmitted	Total number of DMM PDU frames sent to the remote MEP in this measurement session.
Valid packets received	Total number of DMR PDU frames received from the remote MEP in this measurement session.
Average delay	Average two-way frame delay measured in this session.
Average delay variation	Average frame jitter measured in this session.

Table 32: monitor ethernet delay-measurement two-way Output Fields (*continued*)

Output Field Name	Output Field Description
Best case delay	Lowest two-way frame delay measured in this session.
Worst case delay	Highest two-way frame delay measured in this session.

Sample Output

monitor ethernet delay-measurement one-way

```

user@switch> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA

```

monitor ethernet delay-measurement two-way

```

user@switch> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---
Packets transmitted: 10, Valid packets received: 10
Average delay: 103 usec, Average delay variation: 8 usec
Best case delay: 92 usec, Worst case delay: 122 usec

```

monitor ethernet delay-measurement two-way (Invalid DMR Frames Received)

```

user@switch> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a with invalid timestamp(s).
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec

--- Delay measurement statistics ---

```

Packets transmitted: 10, Valid packets received: 9, Invalid packets received: 1
Average delay: 105 usec, Average delay variation: 9 usec
Best case delay: 92 usec, Worst case delay: 122 usec

show oam ethernet connectivity-fault-management delay-statistics

Syntax	<pre>show oam ethernet connectivity-fault-management delay-statistics <count <i>entry-count</i>> <local-mep <i>local-mep-id</i>> maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> <remote-mep <i>remote-mep-id</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 9.5.</p> <p>Command introduced in Junos OS Release 11.4 for EX Series switches.</p>
Description	<p>On MX Series routers with Ethernet interfaces on Dense Port Concentrators (DPCs), display ETH-DM delay statistics.</p> <p>On EX Series switches, display delay measurement results.</p>
Options	<p>count <i>entry-count</i>—(Optional) Number of entries to display from the statistics table. The range of values is 1 through 100. The default value is 100 entries.</p> <p>local-mep <i>local-mep-id</i>—(Optional) Numeric identifier of the local MEP. On MX Series routers, the range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.</p> <p>maintenance-association <i>ma-name</i>—Name of an existing CFM maintenance association.</p> <p>maintenance-domain <i>md-name</i>—Name of an existing connectivity fault management (CFM) maintenance domain.</p> <p>remote-mep <i>remote-mep-id</i>—(Optional) Numeric identifier of the remote MEP. On MX Series routers, the range of values is 1 through 8192. On EX Series switches, the range of values is 1 through 8191.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management statistics • clear oam ethernet connectivity-fault-management delay-statistics on page 392 • show oam ethernet connectivity-fault-management interfaces • show oam ethernet connectivity-fault-management mep-database • show oam ethernet connectivity-fault-management mep-statistics
List of Sample Output	<p>show oam ethernet connectivity-fault-management delay-statistics on page 403</p> <p>show oam ethernet connectivity-fault-management delay-statistics remote-mep on page 403</p>
Output Fields	<p>Table 33 on page 402 lists the output fields for the show oam ethernet connectivity-fault-management delay-statistics command and the show oam ethernet</p>

connectivity-fault-management mep-statistics command. Output fields are listed in the approximate order in which they appear.

Table 33: show oam ethernet connectivity-fault-management delay-statistics and mep-statistics Output Fields

Output Field Name	Field Description
MEP identifier	Maintenance association end point (MEP) numeric identifier.
MAC address	Unicast MAC address configured for the MEP.
Remote MEP count	Number of remote MEPs (unless you specify the remote-mep option).
Remote MEP identifier	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Index	Index number that corresponds to the ETH-DM entry in the CFM database.
One-way delay (usec)	<p>For a one-way ETH-DM session, the frame delay time, in microseconds, measured at the receiver MEP.</p> <p>For a detailed description of one-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
Two-way delay (usec)	<p>For a two-way ETH-DM session, the frame delay time, in microseconds, measured at the initiator MEP.</p> <p>For a detailed description of two-way Ethernet frame delay measurement, see the <i>ITU-T Y.1731 Ethernet Service OAM</i> topics in the <i>Junos OS Network Interfaces Library for Routing Devices</i>.</p>
Average one-way delay	Average one-way frame delay for the statistics displayed.
Average one-way delay variation	Average one-way “frame jitter” for the statistics displayed.
Best-case one-way delay	Lowest one-way frame delay for the statistics displayed.
Worst-case one-way delay	Highest one-way frame delay for the statistics displayed.
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way “frame jitter” for the statistics displayed.
Best-case two-way delay	Lowest two-way frame delay for the statistics displayed.
Worst-case two-way delay	Highest two-way frame delay calculated in this session.

Sample Output

show oam ethernet connectivity-fault-
management
delay-statistics

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6
MEP identifier: 100, MAC address: 00:05:85:73:7b:39
Remote MEP count: 2
Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259      519
    2      273      550
    3      287      571
    4      299      610
    5      313      650
Average one-way delay      : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay    : 259 usec
Worst case one-way delay   : 313 usec
Average two-way delay      : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay    : 519 usec
Worst case two-way delay   : 650 usec

Remote MEP identifier: 102
Remote MAC address: 00:04:55:63:39:5a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      29      58
    2      23      59
    3      27      56
    4      29      62
    5      33      68
Average one-way delay      : 28 usec
Average one-way delay variation: 3 usec
Best case one-way delay    : 23 usec
Worst case one-way delay   : 33 usec
Average two-way delay      : 60 usec
Average two-way delay variation: 3 usec
Best case two-way delay    : 56 usec
Worst case two-way delay   : 68 usec
```

show oam ethernet connectivity-fault-
management delay-statistics remote-mep

```
user@switch> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md6 maintenance-association ma6 remote-mep 101
MEP identifier: 100, MAC address: 00:05:85:73:7b:39

Remote MEP identifier: 101
Remote MAC address: 00:05:85:73:39:4a
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)      (usec)
    1      259      519
```

2	273	550
3	287	571
4	299	610
5	313	650

Average one-way delay : 286 usec
Average one-way delay variation: 62 usec
Best case one-way delay : 259 usec
Worst case one-way delay : 313 usec
Average two-way delay : 580 usec
Average two-way delay variation: 26 usec
Best case two-way delay : 519 usec
Worst case two-way delay : 650 usec

show oam ethernet connectivity-fault-management forwarding-state

Syntax	show oam ethernet connectivity-fault-management forwarding-state interface <i>interface-name</i> <brief detail extensive>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.
Options	interface <i>interface-name</i> —Display forwarding state information for the specified Ethernet interface only. brief detail extensive —(Optional) Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management statistics on page 395 • show oam ethernet connectivity-fault-management path-database on page 415 • show oam ethernet connectivity-fault-management mip on page 424
List of Sample Output	show oam ethernet connectivity-fault-management forwarding-state on page 406 show oam ethernet connectivity-fault-management forwarding-state interface on page 406 show oam ethernet connectivity-fault-management forwarding-state interface detail on page 407 show oam ethernet connectivity-fault-management forwarding-state interface interface-name on page 408
Output Fields	Table 34 on page 405 lists the output fields for the show oam ethernet connectivity-fault-management forwarding-state command. Output fields are listed in the approximate order in which they appear.

Table 34: show oam ethernet connectivity-fault-management forwarding-state Output Fields

Field Name	Field Description	Level of Output
Interface name	Interface identifier.	All levels
Filter action	Filter action for messages at the level.	All levels
Nexthop type	Next-hop type.	All levels
Nexthop index	Next-hop index number.	brief
Level	Maintenance domain (MD) level.	detail

Table 34: show oam ethernet connectivity-fault-management forwarding-state Output Fields (*continued*)

Field Name	Field Description	Level of Output
Direction	MEP direction configured.	none
CEs	Number of customer edge (CE) interfaces.	All levels

Sample Output

show oam ethernet
connectivity-fault-
management forwarding-
state

```
user@host> show oam ethernet connectivity-fault-management forwarding-state
CEs: 3
```

Maintenance domain forwarding state:

Level	Direction	Filter action	Nexthop type	Nexthop index
0		Drop	none	
1		Drop	none	
2		Drop	none	
3		Drop	none	
4		Drop	none	
5		Drop	none	
6		Drop	none	
7		Drop	none	

show oam ethernet
connectivity-fault-
management forwarding-
state interface

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-3/0/0.0
```

Maintenance domain forwarding state:

Level	Direction	Filter action	Nexthop type	Nexthop index
0		Drop	none	
1		Drop	none	
2		Drop	none	
3		Drop	none	
4		Drop	none	
5		Drop	none	
6		Drop	none	
7	down	Receive	none	

```
Interface name: xe-0/0/0.0
```

```
Instance name: __+bd1__
```

Maintenance domain forwarding state:

Level	Direction	Filter action	Nexthop type	Nexthop index
0		Drop	none	

1		Drop	none
2		Drop	none
3		Drop	none
4		Drop	none
5		Drop	none
6		Drop	none
7	down	Receive	none

show oam ethernet
connectivity-fault-
management forwarding-
state interface detail

user@host> show oam ethernet connectivity-fault-management forwarding-state interface
detail

Interface name: ge-3/0/0.0

Level: 0
Filter action: Drop
Nexthop type: none

Level: 1
Filter action: Drop
Nexthop type: none

Level: 2
Filter action: Drop
Nexthop type: none

Level: 3
Filter action: Drop
Nexthop type: none

Level: 4
Filter action: Drop
Nexthop type: none

Level: 5
Filter action: Drop
Nexthop type: none

Level: 6
Filter action: Drop
Nexthop type: none

Level: 7
Direction: down
Filter action: Receive
Nexthop type: none

Interface name: xe-0/0/0.0

Level: 0
Filter action: Drop
Nexthop type: none

Level: 1
Filter action: Drop
Nexthop type: none

...

show oam ethernet
connectivity-fault-
management forwarding-
state interface
interface-name

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface  
interface-name ge-3/0/0.0  
Interface name: ge-3/0/0.0
```

Maintenance domain forwarding state:

Level	Direction	Filter action	Nexthop type	Nexthop index
0		Drop	none	
1		Drop	none	
2		Drop	none	
3		Drop	none	
4		Drop	none	
5		Drop	none	
6		Drop	none	
7	down	Receive	none	

show oam ethernet connectivity-fault-management interfaces

Syntax	show oam ethernet connectivity-fault-management interfaces <ethernet-interface-name> <level md-level> <brief detail extensive>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.
Options	brief detail extensive —(Optional) Display the specified level of output. ethernet-interface-name —(Optional) Display CFM information only for CFM entities attached to the specified Ethernet interface. level md-level —(Optional) Display CFM information for CFM identities enclosed within a maintenance domain of the specified level.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management statistics on page 395 • show oam ethernet connectivity-fault-management path-database on page 415 • show oam ethernet connectivity-fault-management mip on page 424
List of Sample Output	show oam ethernet connectivity-fault-management interfaces on page 412 show oam ethernet connectivity-fault-management interfaces detail on page 412 show oam ethernet connectivity-fault-management interfaces extensive on page 413 show oam ethernet connectivity-fault-management interfaces level on page 414 show oam ethernet connectivity-fault-management interfaces (Trunk Interfaces) on page 414
Output Fields	Table 35 on page 409 lists the output fields for the show oam ethernet connectivity-fault-management interfaces command. Output fields are listed in the approximate order in which they appear.

Table 35: show oam ethernet connectivity-fault-management interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Interface identifier.	All levels
Interface status	Local interface status.	All levels
Link status	Local link status. Up, down, or oam-down.	All levels
Maintenance domain name	Maintenance domain name.	detail extensive

Table 35: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Format (Maintenance domain)	Maintenance domain name format configured.	detail extensive
Level	Maintenance domain level configured.	All levels
Maintenance association name	Maintenance association name.	detail extensive
Format (Maintenance association)	Maintenance association name format configured.	detail extensive
Continuity-check status	Continuity-check status.	detail extensive
Interval	Continuity-check message interval.	detail extensive
Loss-threshold	Lost continuity-check message threshold.	detail extensive
MEP identifier	Maintenance association end point (MEP) identifier.	All levels
Neighbours	Number of MEP neighbors.	All levels
Direction	MEP direction configured.	detail extensive
MAC address	MAC address configured for the MEP.	detail extensive
MEP status	Indicates the status of the Connectivity Fault Management (CFM) protocol running on the MEP: Running , inactive , disabled , or unsupported .	detail extensive
Remote MEP not receiving CCM	Whether the remote MEP is not receiving connectivity check messages (CCMs).	detail extensive
Erroneous CCM received	Whether erroneous CCMs have been received.	detail extensive
Cross-connect CCM received	Whether cross-connect CCMs have been received.	detail extensive
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.	detail extensive
CCMs sent	Number of CCMs transmitted.	detail extensive
CCMs received out of sequence	Number of CCMs received out of sequence.	detail extensive

Table 35: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
LBM sent	Number of loopback request messages (LBMs) sent.	detail extensive
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.	detail extensive
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.	detail extensive
LBRs received with corrupted data	Number of LBRs received that were corrupted.	detail extensive
LBRs sent	Number of LBRs transmitted.	detail extensive
LTMs sent	Linktrace messages (LTMs) transmitted.	detail extensive
LTMs received	Linktrace messages received.	detail extensive
LTRs sent	Linktrace responses (LTRs) transmitted.	detail extensive
LTRs received	Linktrace responses received.	detail extensive
Sequence number of next LTM request	Sequence number of next LTM request to be transmitted.	detail extensive
1DMs sent	<p>If the interface is attached to an initiator MEP for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.</p> <p>For all other cases, this field displays 0.</p>	detail extensive
Valid 1DMs received	<p>If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of valid 1DM frames received.</p> <p>For all other cases, this field displays 0.</p>	detail extensive
Invalid 1DMs received	<p>If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of invalid 1DM frames received.</p> <p>For all other cases, this field displays 0.</p>	detail extensive
DMMs sent	<p>If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.</p> <p>For all other cases, this field displays 0.</p>	detail extensive
DMRs sent	<p>If the interface is attached to a responder MEP for a two-way ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent.</p> <p>For all other cases, this field displays 0.</p>	detail extensive

Table 35: show oam ethernet connectivity-fault-management interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
Valid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays 0.	detail extensive
Invalid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays 0.	detail extensive
Remote MEP count	Number of remote MEPs.	extensive
Identifier (remote MEP)	MEP identifier of the remote MEP.	extensive
MAC address (remote MEP)	MAC address of the remote MEP.	extensive
State (remote MEP)	State of the remote MEP.	extensive
Interface (remote MEP)	Interface of the remote MEP.	extensive

Sample Output

show oam ethernet connectivity-fault-management interfaces

```

user@host> show oam ethernet connectivity-fault-management interfaces
Interface      Link      Status      Level      MEP      Neighbours
               Identifier
ge-1/1/0.0     Up        Active      0          2        1
ge-1/1/0.1     Up        Active      0          2        1
ge-1/1/0.10    Up        Active      0          2        1
ge-1/1/0.100   Up        Active      0          2        1
ge-1/1/0.101   Up        Active      0          2        1
ge-1/1/0.102   Up        Active      0          2        1
ge-1/1/0.103   Up        Active      0          2        1
ge-1/1/0.104   Up        Active      0          2        1
ge-1/1/0.105   Up        Active      0          2        1
ge-1/1/0.106   Up        Active      0          2        1

```

...

show oam ethernet connectivity-fault-management interfaces detail

```

user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string

```

```

Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : yes
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : yes
Statistics:
  CCMs sent                                  : 76
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                             : 0
  LTRs sent                                  : 0
  LTRs received                             : 0
  Sequence number of next LTM request        : 0
  1DMs sent                                  : 0
  Valid 1DMs received                       : 0
  Invalid 1DMs received                     : 0
  DMMs sent                                  : 0
  DMRs sent                                  : 0
  Valid DMRs received                       : 0
  Invalid DMRs received                     : 0
Remote MEP count: 2
  Identifier  MAC address  State  Interface
  2001       00:90:69:0b:7f:71  ok    ge-5/2/9.0
  4001       00:90:69:0b:09:c5  ok    ge-5/2/9.0

```

show oam ethernet connectivity-fault-management interfaces extensive

```

user@host> show oam ethernet connectivity-fault-management interfaces extensive
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
Maintenance domain name: md0, Format: string, Level: 5
Maintenance association name: ma1, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : yes
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : yes
Statistics:
  CCMs sent                                  : 76
  CCMs received out of sequence              : 0
  LBMs sent                                  : 0
  Valid in-order LBRs received               : 0
  Valid out-of-order LBRs received           : 0
  LBRs received with corrupted data          : 0
  LBRs sent                                  : 0
  LTMs sent                                  : 0
  LTMs received                             : 0
  LTRs sent                                  : 0
  LTRs received                             : 0

```

```

Sequence number of next LTM request      : 0
1DMs sent                               : 0
Valid 1DMs received                      : 0
Invalid 1DMs received                   : 0
DMMs sent                               : 0
DMRs sent                               : 0
Valid DMRs received                     : 0
Invalid DMRs received                   : 0
Remote MEP count: 2
Identifier    MAC address      State   Interface
2001         00:90:69:0b:7f:71  ok     ge-5/2/9.0
4001         00:90:69:0b:09:c5  ok     ge-5/2/9.0

```

show oam ethernet connectivity-fault-management interfaces level

```

user@host> show oam ethernet connectivity-fault-management interfaces level 7
Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-3/0/0.0     Up        Active      7          201      0
xe-0/0/0.0     Up        Active      7          203      1

```

show oam ethernet connectivity-fault-management interfaces (Trunk Interfaces)

```

user@host> show oam ethernet connectivity-fault-management interfaces

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-4/0/1.0, vln 100      Up        Active      5          100      0
ge-10/3/10.4091, vln 4091 Down      Inactive    4          400      0
ge-4/0/0.0           Up        Active      6          200      0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-4/0/0.0     Up        Active      6          200      0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vln 100

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-4/0/1.0, vln 100      Up        Active      5          100      0

user@host> show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091
vln 4091

Interface      Link      Status      Level      MEP      Neighbours
Identifier
ge-10/3/10.4091, vln 4091 Down      Inactive    4          400      0

```

show oam ethernet connectivity-fault-management path-database

Syntax	show oam ethernet connectivity-fault-management path-database <i>host</i> maintenance-association <i>ma-name</i> maintenance-domain <i>md-name</i> <i>mac-address</i>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.
Options	<p>mac-address—Display connectivity fault management path database information for the specified MAC address of the remote host.</p> <p>maintenance-association <i>ma-name</i>—Display connectivity fault management path database information for the specified maintenance association.</p> <p>maintenance-domain <i>md-name</i>—Display connectivity fault management path database information for the specified maintenance domain.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management statistics on page 395 • show oam ethernet connectivity-fault-management interfaces on page 409 • show oam ethernet connectivity-fault-management mip on page 424
List of Sample Output	show oam ethernet connectivity-fault-management path-database on page 416 show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 416
Output Fields	Table 36 on page 415 lists the output fields for the show oam ethernet connectivity-fault-management path-database command. Output fields are listed in the approximate order in which they appear.

Table 36: show oam ethernet connectivity-fault-management linktrace path-database Output Fields

Field Name	Field Description
Linktrace to	MAC address of the 802.1ag node to which the linktrace message is targeted.
Interface	Interface used by the local MEP to send the linktrace message (LTM).
Maintenance Domain	Maintenance domain identifier specified in the traceroute command.
Maintenance Association	Maintenance association identifier specified in the traceroute command.
Level	Maintenance domain level configured for the maintenance domain.

Table 36: show oam ethernet connectivity-fault-management linktrace path-database Output Fields (continued)

Field Name	Field Description
Local Mep	MEP identifier of the local MEP originating the linktrace.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message (LTM). The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag maintenance intermediate point (MIP) that is forwarding the LTM.
Next hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming linktrace responses (LTR), with a previously sent LTM.

Sample Output

show oam ethernet connectivity-fault-management path-database

```

user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain MD1 maintenance-association MA1 00:01:02:03:04:05
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
Maintenance Domain: MD1, Level: 7
Maintenance Association: MA1, Local Mep: 1

Hop      TTL      Source MAC address      Next hop MAC address
Transaction Identifier:100001
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc
3         61      00:00:cc:cc:cc:cc      00:01:02:03:04:05
4         60      00:01:02:03:04:05      00:00:00:00:00:00

```

show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands)

```

user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain MD2 maintenance-association MA2 00:06:07:08:09:0A
Linktrace to 00:06:07:08:09:0A, Interface : ge-5/0/1.0
Maintenance Domain: MD2, Level: 6
Maintenance Association: MA2, Local Mep: 10

Hop      TTL      Source MAC address      Next hop MAC address
Transaction Identifier:100002
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc
3         61      00:00:cc:cc:cc:cc      00:06:07:08:09:0A
4         60      00:06:07:08:09:0A      00:00:00:00:00:00
Transaction Identifier:100003
1         63      00:00:aa:aa:aa:aa      00:00:bb:bb:bb:bb
2         62      00:00:bb:bb:bb:bb      00:00:cc:cc:cc:cc

```

3	61	00:00:cc:cc:cc:cc	00:06:07:08:09:0A
4	60	00:06:07:08:09:0A	00:00:00:00:00:00

show oam ethernet connectivity-fault-management mep-database

Syntax	<pre>show oam ethernet connectivity-fault-management mep-database maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> <local-mep <i>local-mep-id</i>> <remote-mep <i>remote-mep-id</i>></pre>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.
Options	<p>maintenance-association <i>ma-name</i>—Display connectivity fault management information for the specified maintenance association.</p> <p>maintenance-domain <i>domain-name</i>—Display connectivity fault management information for the specified maintenance domain.</p> <p>local-mep <i>local-mep-id</i>—(Optional) Display connectivity fault management information for the specified local MEP only.</p> <p>remote-mep <i>remote-mep-id</i>—(Optional) Display connectivity fault management information for the specified remote MEP only.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • clear oam ethernet connectivity-fault-management statistics on page 395 • show oam ethernet connectivity-fault-management interfaces on page 409 • show oam ethernet connectivity-fault-management mip on page 424
List of Sample Output	<p>show oam ethernet connectivity-fault-management mep-database on page 422</p> <p>show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 422</p> <p>show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 422</p>
Output Fields	Table 37 on page 418 lists the output fields for the show oam ethernet connectivity-fault-management mep-database command. Output fields are listed in the approximate order in which they appear.

Table 37: show oam ethernet connectivity-fault-management mep-database Output Fields

Field Name	Field Description
Maintenance domain name	Maintenance domain name.

Table 37: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

Field Name	Field Description
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.
Format (Maintenance association)	Maintenance association name format configured.
Continuity-check status	Continuity-check status.
Interval	Continuity-check message interval.
MEP identifier	Maintenance association end point (MEP) identifier.
Direction	MEP direction configured.
MAC address	MAC address configured for the MEP.
Auto-discovery	Whether automatic discovery is enabled or disabled.
Priority	Priority used for CCMs and linktrace messages transmitted by the MEP.
Interface name	Interface identifier.
Interface status	Local interface status.
Link status	Local link status.
Remote MEP not receiving CCM	Whether the remote MEP is not receiving CCMs.
Erroneous CCM received	Whether erroneous CCMs have been received.
Cross-connect CCM received	Whether cross-connect CCMs have been received.
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.

Table 37: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

Field Name	Field Description
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.
LBRs received with corrupted data	Number of LBRs received that were corrupted.
LBRs sent	Number of LBRs transmitted.
LTMs sent	Linktrace messages (LTMs) transmitted.
LTMs received	Linktrace messages received.
LTRs sent	Linktrace responses (LTRs) transmitted.
LTRs received	Linktrace responses received.
Sequence number of next LTM request	Sequence number of the next linktrace message request to be transmitted.
1DMs sent	<p>If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session.</p> <p>For all other cases, this field displays 0.</p>
Valid 1DMs received	<p>If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received.</p> <p>For all other cases, this field displays 0.</p>
Invalid 1DMs received	<p>If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received.</p> <p>For all other cases, this field displays 0.</p>
DMMs sent	<p>If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session.</p> <p>For all other cases, this field displays 0.</p>
DMRs sent	<p>If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent.</p> <p>For all other cases, this field displays 0.</p>

Table 37: show oam ethernet connectivity-fault-management mep-database Output Fields (*continued*)

Field Name	Field Description
Valid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays 0.
Invalid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays 0.
Remote MEP identifier	MEP identifier of the remote MEP.
State (remote MEP)	State of the remote MEP: idle , start , ok , or failed .
MAC address	MAC address of the remote MEP.
Type	Whether the remote MEP MAC address was learned using automatic discovery or configured.
Interface	Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.
Last flapped	Date, time, and how long ago the remote MEP interface went from down to up. The format is Last flapped: year-month-day hours:minutes:seconds timezone (hours:minutes:seconds ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .
Remote defect indication	Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.
Port status TLV	<ul style="list-style-type: none"> In the Maintenance domain section, displays the last transmitted port status TLV value. In the Remote MEP section, displays the last value of port status TLV received from the remote MEP. <p>In the Action profile section, displays, the last occurred event port-status-tlv blocked event. This event occurred due to the reception of blocked value in the port status TLV from remote MEP.</p>
Interface status TLV	<ul style="list-style-type: none"> In the Maintenance domain section, displays the last transmitted interface status TLV value. In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP. <p>In the Action profile section, if displays, the last occurred event interface-status-tlv event (either lower-layer-down or down). This event occurred due to the reception of either lower or down value in the interface status TLV from remote MEP.</p>
Action profile	Name of the action profile occurrence associated with a remote MEP.
Last event	When an action profile occurs, displays the last event that triggered it.
Last event cleared	When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.
Action	Action taken and the corresponding time of the action occurrence.

Sample Output

show oam ethernet connectivity-fault-management mep-database

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                      : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                       : no
Statistics:
  CCMs sent                                  : 1476
  CCMs received out of sequence              : 0
  LBMs sent                                  : 85
Remote MEP count: 1
Identifier   MAC address      State   Interface
  100       00:19:e2:b2:81:4b   ok     vt-0/1/10.1049088

```

show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200
remote-mep 100
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
Maintenance association name: vpls-vlan200, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

Remote MEP identifier: 100, State: ok
MAC address: 00:19:e2:b2:81:4b, Type: Learned
Interface: vt-0/1/10.1049088
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none

```

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none
Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

```

Remote MEP identifier: 200, State: ok
MAC address: 00:05:85:73:96:1f, Type: Configured
Interface: ge-1/0/8.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: lower-layer-down
Action profile: juniper
 Last event: Interface-status-tlv lower-layer-down
 Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)

show oam ethernet connectivity-fault-management mip

Syntax	show oam ethernet connectivity-fault-management mip <qualifier>
Release Information	Command introduced in Junos OS Release 10.2 for EX Series switches.
Description	Display all the maintenance association intermediate points (MIPs) created in the system. Specify qualifiers to display specific MIPs.
Options	<i>qualifier</i> —(Optional) Display the specified MIP.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show oam ethernet connectivity-fault-management interfaces on page 409 • show oam ethernet connectivity-fault-management path-database on page 415
List of Sample Output	show oam ethernet connectivity-fault-management mip on page 424
Output Fields	Table 38 on page 424 lists the output fields for the show oam ethernet connectivity-fault-management mip command. Output fields are listed in the approximate order in which they appear.

Table 38: show oam ethernet connectivity-fault-management mip Output Fields

Field Name	Field Description
MIP information for instance	Header for the MIP information showing the MIP name.
Interface	Interface type-dpc/pic/port.unit-number.
Level	MIP level configured.

Sample Output

show oam ethernet connectivity-fault-management mip

```

user@host> show oam ethernet connectivity-fault-management mip
MIP information for  __mip_name__

MIP information for  default-switch bd1

    Interface      Level
    ge-3/0/0.0     7
    ge-3/0/1.0     6
    ge-3/0/3.0     6

```

show oam ethernet connectivity-fault-management sla-iterator-statistics

Syntax	<pre>show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> sla-iterator <i>sla-iterator</i> <local-mep <i>local-mep-id</i>> <remote-mep <i>remote-mep-id</i>></pre>
Release Information	<p>Command introduced in Junos OS Release 11.4 for EX Series switches.</p> <p>Command introduced in Junos OS Release 9.6.</p> <p>Command introduced in Junos OS Release 12.2 for ACX Series routers.</p> <p>Command introduced in Junos OS Release 13.2 for MX Series routers (not on MPC3E Hyperion cards).</p>
Description	Display the Ethernet Operation, Administration, and Maintenance (OAM) service-level agreement (SLA) iterator statistics.
Options	<p>maintenance-domain <i>md-name</i>—Name of an existing connectivity fault management (CFM) maintenance domain.</p> <p>maintenance-association <i>ma-name</i>—Name of an existing CFM maintenance association.</p> <p>sla-iterator <i>sla-iterator</i>— Name of the iterator profile.</p> <p>local-mep <i>local-mep-id</i>—(Optional) Numeric identifier of the local MEP. The range of values is 1 through 8191.</p> <p>remote-mep <i>remote-mep-id</i>—(Optional) Numeric identifier of the remote MEP. The range of values is 1 through 8192.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring an Iterator Profile on a Switch (CLI Procedure) on page 102 • clear oam ethernet connectivity-fault-management sla-iterator-statistics on page 394
List of Sample Output	<p>show oam ethernet connectivity-fault-management sla-iterator-statistics on page 428</p> <p>show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers) on page 428</p>
Output Fields	Table 39 on page 425 lists the output fields for the show oam ethernet connectivity-fault-management sla-iterator-statistics command. Output fields are listed in the approximate order in which they appear.

Table 39: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields

Output Field Name	Output Field Description
Maintenance domain	Name of the maintenance domain.

Table 39: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

Output Field Name	Output Field Description
Level	Level of the maintenance domain level configured.
Maintenance association	Name of the maintenance association.
Local MEP id	Numeric identifier of the local MEP.
Remote MEP id	Numeric identifier of the remote MEP.
Remote MAC address	Unicast MAC address of the remote MEP.
Iterator name	Name of iterator.
Iterator Id	Numeric identifier of the iterator.
Iterator cycle time	Number of cycles (in milliseconds) taken between back-to-back transmission of SLA frames for this connection
Iteration period	Maximum number of cycles per iteration
Iterator status	Current status of iterator whether running or stopped.
Infinite iterations	Status of iteration as infinite or finite.
Counter reset time	Date and time when the counter was reset.
Reset reason	Reason to reset counter.
Delay weight	Calculation weight of delay.
Delay variation weight	Calculation weight of delay variation.
DMM sent	Delay measurement message (DMM) PDU frames sent to the peer MEP in this session.
DMM skipped for threshold hit	Number of DMM frames sent to the peer MEP in this session skipped during threshold hit.
DMM skipped for threshold hit window	Number of DMM frames sent to the peer MEP in this session skipped during the last threshold hit window.
DMR received	Number of delay measurement reply (DMR) frames received.
DMR out of sequence	Total number of DMR out of sequence packets received.
DMR received with invalid time stamps	Total number of DMR frames received with invalid timestamps.

Table 39: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

Output Field Name	Output Field Description
Average two-way delay	Average two-way frame delay for the statistics displayed.
Average two-way delay variation	Average two-way "frame jitter" for the statistics displayed.
Average one-way forward delay variation	Average one-way forward delay variation for the statistics displayed in microseconds.
Average one-way backward delay variation	Average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average two-way delay	Weighted average two-way delay for the statistics displayed in microseconds.
Weighted average two-way delay variation	Weighted average two-way delay variation for the statistics displayed in microseconds.
Weighted average one-way backward delay variation	Weighted average one-way backward delay variation for the statistics displayed in microseconds.
Weighted average one-way forward delay variation	Weighted average one-way forward delay variation for the statistics displayed in microseconds.
SLM packets sent	Total number of synthetic loss message (SLM) PDU frames sent from the source MEP to the remote MEP during this ETH-SLM session.
SLM packets received	Total number of synthetic loss message (SLM) PDU frames that the remote MEP received from the source MEP during this ETH-SLM session.
SLR packets sent	Total number of synthetic loss reply (SLR) PDU frames that the remote MEP sent to the source MEP during this measurement session.
SLR packets received	Total number of synthetic loss reply (SLR) PDU frames that the source MEP received from the remote MEP during this measurement session.
Local TXFC1 value	Number of synthetic frames transmitted to the peer MEP for a test ID. A test ID is used to distinguish each synthetic loss measurement because multiple measurements can be simultaneously activated also on a given CoS and MEP pair. It must be unique at least within the context of any SLM for the MEG and initiating MEP.
Local RXFC1 value	Number of synthetic frames received from the peer MEP for a test ID. The MEP generates a unique Test ID for the session, adds the source MEP ID, and initializes the local counters for the session before SLM initiation. For each SLM PDU transmitted for the session (test ID), the local counter TXFC1 is sent in the packet.
Last Received SLR frame TXFCf(tc)	Value of the local counter TxFC1 at the time of SLM frame transmission.

Table 39: show oam ethernet connectivity-fault-management sla-iterator-statistics Output Fields (*continued*)

Output Field Name	Output Field Description
Last Received SLR frame TXFCb(t	Value of the local counter RxFCI at the time of SLR frame transmission.
Frame loss (near-end)	Count of frame loss associated with ingress data frames.
Frame loss (far-end)	Count of frame loss associated with egress data frames.

Sample Output

show oam ethernet connectivity-fault-management sla-iterator-statistics

```

user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics
sla-iterator i1 maintenance-domain default-1 maintenance-association ma1 local-mep 1
remote-mep 2
Iterator statistics:
Maintenance domain: md6, Level: 6
Maintenance association: ma6, Local MEP id: 1000
Remote MEP id: 103, Remote MAC address: 00:90:69:0a:43:92
Iterator name: i1, Iterator Id: 1
Iterator cycle time: 10ms, Iteration period: 1 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2010-03-19 20:42:39 PDT (2d 18:24 ago)
Reset reason: Adjacency flap

Iterator delay measurement statistics:
Delay weight: 1, Delay variation weight: 1
DMM sent : 23898520
DMM skipped for threshold hit : 11000
DMM skipped for threshold hit window : 0
DMR received : 23851165
DMR out of sequence : 1142
DMR received with invalid time stamps : 36540
Average two-way delay : 129 usec
Average two-way delay variation : 15 usec
Average one-way forward delay variation : 22 usec
Average one-way backward delay variation : 22 usec
Weighted average two-way delay : 134 usec
Weighted average two-way delay variation : 8 usec
Weighted average one-way forward delay variation : 6 usec
Weighted average one-way backward delay variation : 2 usec

```

Sample Output

show oam ethernet connectivity-fault-management sla-iterator-statistics (MX Series routers)

```

user@switch> show oam ethernet connectivity-fault-management sla-iterator-statistics
maintenance-domain md0 maintenance-association mau local-mep 4 remote-mep 3 sla-iterator
lm
Iterator statistics:
Maintenance domain: 2, Level: 2
Maintenance association: W-160432000-001, Local MEP id: 2
Remote MEP id: 1, Remote MAC address: 00:90:69:0a:43:39

```

```
Iterator name: iter1, Iterator Id: 1
Iterator cycle time: 100ms, Iteration period: 10 cycles
Iterator status: running, Infinite iterations: true
Counter reset time: 2012-09-25 02:15:31 PDT (00:00:45 ago)
Reset reason: Adjacency flap
Iterator loss measurement statistics:
  LMM sent : 444
  LMM skipped for threshold hit : 0
  LMM skipped for threshold hit window: 0
  LMR received : 444
  LMR out of sequence : 0
  LMR forwarding-class mismatch : 0
Accumulated transmit statistics:
  Near-end (CIR) : 0
  Far-end (CIR) : 0
  Near-end (EIR) : 0
  Far-end (EIR) : 0
Accumulated receive statistics:
  Near-end (CIR) : 0
  Far-end (CIR) : 0
  Near-end (EIR) : 0
  Far-end (EIR) : 0
Accumulated loss statistics:
  Near-end loss (CIR) : 0
  Near-end loss-ratio (CIR) : 0 (0.00000%)
  Far-end loss (CIR) : 0
  Far-end loss-ratio (CIR) : 0 (0.00000%)
  Near-end loss (EIR) : 0
  Near-end loss-ratio (EIR) : 0 (0.00000%)
  Far-end loss (EIR) : 0
  Far-end loss-ratio (EIR) : 0 (0.00000%)
Last loss measurement statistics:
  Near-end (CIR) : 0
  Far-end (CIR) : 0
  Near-end (EIR) : 0
  Far-end (EIR) :
```


CHAPTER 22

Operational Commands: Ethernet OAM Link Fault Management

- `show oam ethernet link-fault-management`

show oam ethernet link-fault-management

Syntax	show oam ethernet link-fault-management <brief detail> <interface-name>
Release Information	Command introduced in Junos OS Release 9.4 for EX Series switches.
Description	Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.
Options	brief detail —(Optional) Display the specified level of output. interface-name —(Optional) Display link fault management information for the specified Ethernet interface only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring Ethernet OAM Link Fault Management on EX Series Switches on page 72 • Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 97
List of Sample Output	show oam ethernet link-fault-management brief on page 436 show oam ethernet link-fault-management detail on page 436
Output Fields	Table 40 on page 432 lists the output fields for the show oam ethernet link-fault-management command. Output fields are listed in the approximate order in which they appear.

Table 40: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Indicates the status of the established link. <ul style="list-style-type: none"> • Fail—A link fault condition exists. • Running—A link fault condition does not exist. 	All levels
Discovery state	State of the discovery mechanism: <ul style="list-style-type: none"> • Passive Wait • Send Any • Send Local Remote • Send Local Remote Ok 	All levels
Peer address	Address of the OAM peer.	All levels

Table 40: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flags	Information about the interface. <ul style="list-style-type: none"> • Remote-Stable—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. False indicates that remote DTE has either not seen or is unsatisfied with local state information. True indicates that remote DTE has seen and is satisfied with local state information. • Local-Stable—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. False indicates that local DTE either has not seen or is unsatisfied with remote state information. True indicates that local DTE has seen and is satisfied with remote state information. • Remote-State-Valid—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. False indicates that OAM client has not seen remote state information. True indicates that the OAM client has seen remote state information. 	All levels
Remote loopback status	Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
Remote entity information	Remote entity information. <ul style="list-style-type: none"> • Remote MUX action—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • Remote parser action—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • Discovery mode—Indicates whether discovery mode is active or inactive. • Unidirectional mode—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes. • Remote loopback mode—Indicates whether remote loopback is supported or not supported. • Link events—Indicates whether interpreting link events is supported or not supported on the remote peer. • Variable requests—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer. 	All levels
OAM Receive Statistics		
Information	The number of information PDUs received.	detail
Event	The number of loopback control PDUs received.	detail
Variable request	The number of variable request PDUs received.	detail
Variable response	The number of variable response PDUs received.	detail
Loopback control	The number of loopback control PDUs received.	detail

Table 40: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
Organization specific	The number of vendor organization specific PDUs received.	detail
OAM Transmit Statistics		
Information	The number of information PDUs transmitted.	detail
Event	The number of event notification PDUs transmitted.	detail
Variable request	The number of variable request PDUs transmitted.	detail
Variable response	The number of variable response PDUs transmitted.	detail
Loopback control	The number of loopback control PDUs transmitted.	detail
Organization specific	The number of vendor organization specific PDUs transmitted.	detail
OAM Received Symbol Error Event information		
Events	The number of symbol error event TLVs that have been received after the OAM sublayer was reset.	detail
Window	The symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	detail
Threshold	The number of errored symbols in the period required for the event to be generated.	detail
Errors in period	The number of symbol errors in the period reported in the received event PDU.	detail
Total errors	The number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	detail
OAM Received Frame Error Event Information		
Events	The number of errored frame event TLVs that have been received after the OAM sublayer was reset.	detail
Window	The duration of the window in terms of the number of 100 ms period intervals.	detail
Threshold	The number of detected errored frames required for the event to be generated.	detail
Errors in period	The number of detected errored frames in the period.	detail

Table 40: show oam ethernet link-fault-management Output Fields (*continued*)

Field Name	Field Description	Level of Output
Total errors	The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	detail
OAM Received Frame Period Error Event Information		
Events	The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	detail
Window	The duration of the frame seconds window.	detail
Threshold	The number of frame seconds errors in the period.	detail
Errors in period	The number of frame seconds errors in the period.	detail
Total errors	The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	detail
OAM Transmitted Symbol Error Event Information		
Events	The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	detail
Window	The symbol error event window in the transmitted PDU.	detail
Threshold	The number of errored symbols in the period required for the event to be generated.	detail
Errors in period	The number of symbol errors in the period reported in the transmitted event PDU.	detail
Total errors	The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	detail
OAM Transmitted Frame Error Event Information		
Events	The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	detail
Window	The duration of the window in terms of the number of 100 ms period intervals.	detail
Threshold	The number of detected errored frames required for the event to be generated.	detail
Errors in period	The number of detected errored frames in the period.	detail
Total errors	The number of errored frames that have been detected after the OAM sublayer was reset.	detail

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:72:2c:83
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote loopback status: Disabled on local port, Enabled on peer port
Remote entity information:
  Remote MUX action: discarding, Remote parser action: loopback
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
Interface: ge-0/0/1
Status: Running, Discovery state: Send Any
Peer address: 00:90:69:0a:07:14
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
OAM receive statistics:
  Information: 186365, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM transmit statistics:
  Information: 186347, Event: 0, Variable request: 0, Variable response: 0
  Loopback control: 0, Organization specific: 0
OAM received symbol error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM received frame period error event information:
  Events: 0, Window: 0, Threshold: 0
  Errors in period: 0, Total errors: 0
OAM transmitted symbol error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
OAM transmitted frame error event information:
  Events: 0, Window: 0, Threshold: 1
  Errors in period: 0, Total errors: 0
Remote entity information:
  Remote MUX action: forwarding, Remote parser action: forwarding
  Discovery mode: active, Unidirectional mode: unsupported
  Remote loopback mode: supported, Link events: supported
  Variable requests: unsupported
```

CHAPTER 23

Operational Commands: Uplink Failure Detection

- `show uplink-failure-detection`

show uplink-failure-detection

Syntax	<code>show uplink-failure-detection</code> <code><group group-name></code>
Release Information	Command introduced in Junos OS Release 11.1 for EX Series switches.
Description	Display information about the uplink-failure-detection group, the member interfaces, and their status.
Options	none —Display information about all groups configured for uplink failure detection. group group-name —(Optional) Display information about the specified group only.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 99
List of Sample Output	show uplink-failure-detection on page 438
Output Fields	Table 41 on page 438 lists the output fields for the show uplink-failure-detection command. Output fields are listed in the approximate order in which they appear.

Table 41: show uplink-failure-detection Output Fields

Field Name	Field Description
Group	Name of the group.
Uplink	The uplink interface or interfaces configured as link-to-monitor. NOTE: The asterisk (*) indicates that the link is up.
Downlink	The downlink interface or interfaces configured as link-to-disable. NOTE: The asterisk (*) indicates that the link is up.
Failure Action	Status of uplink failure detection: <ul style="list-style-type: none"> • Active—The switch has detected an uplink failure and has brought the downlink down. • Inactive—The uplink or uplinks are up.

Sample Output

show uplink-failure-detection

```

user@switch> show uplink-failure-detection
Group           : group1
Uplink          : ge-0/0/0*
Downlink        : ge-0/0/1*
Failure Action   : Inactive

```

```
Group          : group2
Uplink         : ge-0/0/3.0
Downlink       : ge-0/0/4.0
Failure Action : Active
```


PART 4

Troubleshooting

- [Troubleshooting Procedures and Commands on page 443](#)

Troubleshooting Procedures and Commands

- [Troubleshooting Port Mirroring Configuration Error Messages on page 443](#)
- [Troubleshooting sFlow on Management Interface on page 444](#)
- `show pfe statistics bridge`

Troubleshooting Port Mirroring Configuration Error Messages

Troubleshooting issues with port mirroring on EX Series switches:

1. [An Analyzer Configuration Returns a “Multiple interfaces cannot be configured as a member of Analyzer output VLAN” Error Message on page 443](#)

An Analyzer Configuration Returns a “Multiple interfaces cannot be configured as a member of Analyzer output VLAN” Error Message

Problem **Description:** In an analyzer configuration, if the VLAN to which mirrored traffic is sent contains more than one member interface, the following error message is displayed in the CLI when you commit the analyzer configuration and the commit fails:

```
Multiple interfaces cannot be configured as a member of Analyzer output VLAN <vlan
name>
```

Solution You must direct the mirrored traffic to a VLAN that has a single member interface. You can do this by completing either of these tasks:

- Reconfigure the existing VLAN to contain a single member interface. You can choose this method if you want to use the existing VLAN.
- Create a new VLAN with a single member interface and associate the VLAN with the analyzer.

To reconfigure the existing VLAN to contain only one member interface:

1. Remove member interfaces from the VLAN repeatedly by using either the **delete vlan** command or the **delete interface** command until the VLAN contains a single member interface:

- [edit]
user@switch# **delete vlan *vlan-id* interface *interface-name***
 - [edit]
user@switch# **delete interface *interface-name* unit 0 family *family-name* vlan member *vlan-id***
2. (Optional) Confirm that the VLAN contains only one interface:
- [edit]
user@switch# **show vlans *vlan-name***
- The output for this command must display only one interface.

To create a new VLAN with a single member interface:

1. Configure a VLAN to carry the mirrored traffic:
[edit]
user@switch# **set vlans *vlan-name***
2. Associate an interface with the VLAN:
[edit]
user@switch# **set interfaces *interface-name* unit *logical-unit-number* family *family-name* vlan members *vlan-name***
3. Associate the VLAN with the analyzer:
[edit ethernet-switching-options]
user@switch# **set analyzer *analyzer-name* output vlan *vlan-name***

**Related
Documentation**

- [Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on EX Series Switches on page 42](#)
- [Configuring Port Mirroring to Analyze Traffic \(CLI Procedure\) on page 78](#)
- [Configuring Port Mirroring to Analyze Traffic \(J-Web Procedure\) on page 81](#)
- [Understanding Port Mirroring on EX Series Switches on page 9](#)

Troubleshooting sFlow on Management Interface

Problem **Description:** On EX Series switches, sFlow does not work with the management interface me0 as the Collector IP, when there is an alternate route to the collector through a network port.

Cause The sFlow packets are sent through the me0 interface only if route lookup to the collector fails at the Packet Forwarding Engine; that is, when there are no network port routes to the collector. Sometimes the static route to the collector through me0 and the default route through network ports match. The sFlow packets are then sent through the network ports and not through the me0 port. This might result in the sFlow packets being sent through network ports on which no collector is present.

Solution Push the collector route-information to the Packet Forwarding Engine by setting a static **receive** route towards the collector. This will ensure that the sFlow packets are sent through the me0 ports.

Set a static **receive** route to the collector for appropriate forwarding of the packet.

1. Assume the collector IP **192.168.18.166** is IP1. Add a static route to IP1, as the best route through me0 to reach the collector.
2. Add an additional static **receive** route, which is the next best route to the collector.

For example:

```
[edit routing-options static]
root# show
```

- Best route to the collector through me0:
192.168.18.166/32 next-hop 192.168.148.11
- Next best route to the collector:
192.168.18.164/30

- Related Documentation**
- [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 17](#)
 - [Configuring sFlow Technology for Network Monitoring \(CLI Procedure\) on page 84](#)

show pfe statistics bridge

Syntax	show pfe statistics bridge <fpc slot>
Release Information	Command introduced in Junos OS Release 12.1 for EX Series switches.
Description	Display information about the number of packets discarded in the ingress pipeline of the Packet Forwarding Engine, packets discarded because of egress filtering or congestion filtering, number of control packets, and general counters for dropped packets. You can use this information to inform troubleshooting investigations.
Options	none —Display bridge counter statistics for all Flexible PIC Concentrator (FPC) slots. fpc slot —(Optional) Display bridge counter statistics for a specific FPC slot.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Monitoring System Process Information</i> • <i>Monitoring Switch Control Traffic</i>
List of Sample Output	show pfe statistics bridge (EX3200 and EX4200 Switches) on page 447 show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis) on page 448 show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis) on page 449 show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card) on page 449
Output Fields	Table 15 on page 320 lists the output fields for the show pfe statistics bridge command. Output fields are listed in the approximate order in which they appear.

Table 42: show pfe statistics bridge Output Fields

Field Name	Field Description
Ingress Counters	<p>Information about ingress counters:</p> <ul style="list-style-type: none"> • Received—Number of packets received by the bridge. • VLAN Filtered—Number of packets discarded because of VLAN filtering. • Security Filtered—Number of packets discarded because of security filtering. • Other Discards—Number of packets dropped by the bridge for reasons other than VLAN or security filtering.

Table 42: show pfe statistics bridge Output Fields (*continued*)

Field Name	Field Description
Egress Counters	<p>Information about egress counters:</p> <ul style="list-style-type: none"> • Unicast—Number of unicast packets transmitted. • Multicast—Number of multicast packets transmitted. • Broadcast—Number of broadcast packets transmitted. • Egress Filtered—Number of egress-filtered packets (regardless of port, priority, or mode). • TailDrop—Number of packets filtered because of egress queue congestion. • Forward Restrict—Number of packets filtered because of egress forward restrictions. • Congestion Filtered—Number of packets filtered because of transmit queue (TxQ) congestion. • Control Packets—Number of control packets (sent to CPU, received from CPU, and sent to analyzer).
Drop Counters	<p>Information about drop counters:</p> <ul style="list-style-type: none"> • Drop Mode—Count mode of the counter. • Drop Counter—Counter value.
General Counters	<p>Information about general counters:</p> <ul style="list-style-type: none"> • Drop Mode—Count mode of the counter. • Drop Counter—Counter value. • Source Not Learnt—Number of source addresses that were not learnt because of internal congestion.
MUX PFE	<p>Information about multiplexer PFE for oversubscribed cards:</p> <ul style="list-style-type: none"> • Drop Mode—Count mode of the counter. • Drop Count—Counter value.

Sample Output

show pfe statistics bridge (EX3200 and EX4200 Switches)

```

user@switch> show pfe statistics bridge
Slot 0

PFE:                0          1          2
-----
---- Ingress Counters ----
Received:            0          52          0
VLAN Filtered:       0          0          0
Security Filtered:   0          0          0
Other Discards:      0          0          0
---- Egress Counters ----
Unicast:             0         104         52
Multicast:            0          0          0
Broadcast:            0          0          0
Egress Filtered:     0          0          0
Congestion Filtered: 0          0          0
Control Packets:     5          0          0
---- General Counters ----
Drop Mode:           0          0          0

```

```

Drop Counter:          34217          36080          6367
Source Not Learnt:      0              0              0

```

show pfe statistics bridge (EX8200 Switches and EX8200 Virtual Chassis)

```
user@switch> show pfe statistics bridge
```

```
Slot 0
```

```
PFE:                  0              1
```

```
----- Ingress Counters -----
```

```

Received:              946              48
VLAN Filtered:         0              0
Security Filtered:     0              0
Other Discards:        0              0

```

```
----- Egress Counters -----
```

```

Unicast:               0              0
Multicast:             0              0
Broadcast:             0              0
Egress Filtered:       0              0
TailDrop:              0              0
Forward Restrict:      0              0
Congestion Filtered:   0              0
Control Packets:       4103           896

```

```
----- Drop Counters -----
```

```

Drop Mode:             0              0
Drop Counter:          12528          2

```

```
Slot 1
```

```
PFE:                  0              1
```

```
----- Ingress Counters -----
```

```

Received:              0              0
VLAN Filtered:         0              0
Security Filtered:     0              0
Other Discards:        0              0

```

```
----- Egress Counters -----
```

```

Unicast:               0              0
Multicast:             0              0
Broadcast:             0              0
Egress Filtered:       0              0
TailDrop:              0              0
Forward Restrict:      0              0
Congestion Filtered:   0              0
Control Packets:       0              0

```

```
----- Drop Counters -----
```

```

Drop Mode:             0              0
Drop Counter:          0              0

```

```
Slot 2
```

```
PFE:                  0              1
```

```
----- Ingress Counters -----
```

```

Received:              0              0
VLAN Filtered:         0              0
Security Filtered:     0              0
Other Filtered:        0              0

```

```
----- Egress Counters -----
```

```

Unicast:               0              0
Multicast:             0              0

```

```

Broadcast:                0          0
Egress Filtered:          0          0
TailDrop:                 0          0
Forward Restrict:         0          0
Congestion Filtered:     0          0
Control Packets:          0          0
---- Drop Counters ----
Drop Mode:                0          0
Drop Counter:             0          0

```

show pfe statistics bridge fpc (EX8200 Switches and EX8200 Virtual Chassis)

```

user@switch> show pfe statistics bridge fpc 2
Slot 2

```

```

PFE:                0          1
-----
---- Ingress Counters ----
Received:           0          0
VLAN Filtered:      0          0
Security Filtered:  0          0
Other Discards:     0          0
---- Egress Counters ----
Unicast:            0          0
Multicast:          0          0
Broadcast:          0          0
Egress Filtered:    0          0
TailDrop:           0          0
Forward Restrict:   0          0
Congestion Filtered: 0          0
Control Packets:    0          0
---- Drop Counters ----
Drop Mode:          0          0
Drop Counter:       0          0

```

show pfe statistics bridge fpc (EX8200-40XS (40-port SFP+) Line Card)

```

user@switch> show pfe statistics bridge fpc 8
Slot 8

```

```

PFE:                0          1          2          3
-----
---- Ingress Counters ----
Received:           0          3          0          0
VLAN Filtered:      0          0          0          0
Security Filtered:  0          0          0          0
Other Discards:     0          1          0          0
---- Egress Counters ----
Unicast:            0          0          0          0
Multicast:          0          0          0          0
Broadcast:          0          0          0          0
Egress Filtered:    0          0          0          0
TailDrop:           0          0          0          0
Forward Restrict:   0          0          0          0
Congestion Filtered: 0          2          0          0
Control Packets:    4          0          0          0
---- Drop Counters ----
Drop Mode:          0          0          0          0
Drop Counter:       0          1          0          0

MUX PFE:            4          5

```

Drop Mode: 0 0
Drop Count: 0 0