Network Management and Monitoring Feature Guide for EX9200 Switches

Release
16.2

Modified: 2016-11-03

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

Network Management and Monitoring Feature Guide for EX9200 Switches
16.2

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 xi
Documentation and Release Notes xi
Supported Platforms xi
Using the Examples in This Manual xi
Merging a Full Example xii
Merging a Snippet xii
Documentation Conventions xiii
Documentation Feedback xv
Requesting Technical Support xv
Self-Help Online Tools and Resources xv
Opening a Case with JTAC xvi

Part 1 Configuring Network Management

Chapter 1 **Configuring Ethernet OAM Link Fault Management 3**

IEEE 802.3ah OAM Link-Fault Management Overview 3
Configuring IEEE 802.3ah OAM Link-Fault Management 5
Enabling IEEE 802.3ah OAM Support 6
Configuring the OAM PDU Interval 6
Configuring the OAM PDU Threshold 7
Configuring an OAM Action Profile 8
Configuring Threshold Values for Fault Events in an Action Profile 10
Applying an Action Profile 10
Setting a Remote Interface into Loopback Mode 11
Monitoring the Loss of Link Adjacency 12
Monitoring Protocol Status 13
Enabling Remote Loopback Support on the Local Interface 15
Configuring Link Discovery 16
Configuring Threshold Values for Local Fault Events on an Interface 17
Disabling the Sending of Link Event TLVs 18
Detecting Remote Faults 18
Specifying the Actions to Be Taken for Link-Fault Management Events 19
Example: Configuring IEEE 802.3ah OAM Support on an Interface 21

Chapter 2 **Configuring sFlow Monitoring Technology 23**

Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch 23
Sampling Mechanism and Architecture of sFlow Technology on EX Series Switches 23
Adaptive Sampling 24

sFlow Agent Address Assignment . . . 25
Configuring sFlow Technology for Network Monitoring (CLI Procedure) . . . 26
Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches . . . 27

Chapter 3 Configuring Uplink Failure Detection . . . 33

Understanding Uplink Failure Detection . . . 33
Uplink Failure Detection Overview . . . 33
Failure Detection Pair . . . 34
Configuring Interfaces for Uplink Failure Detection (CLI Procedure) . . . 35

Chapter 4 Routine Monitoring . . . 37

Monitoring Traffic Through the Router or Switch . . . 37
Displaying Real-Time Statistics About All Interfaces on the Router or Switch . . . 37
Displaying Real-Time Statistics About an Interface on the Router or Switch . . . 38

Part 2 Configuration Statements and Operational Commands

Chapter 5 Configuration Statements: Ethernet OAM Link Fault Management . . . 43

link-fault-management . . . 45
action (OAM) . . . 46
action-profile (Defining for LFM) . . . 47
allow-remote-loopback . . . 48
apply-action-profile . . . 48
ethernet (Protocols OAM) . . . 49
event (LFM) . . . 52
event-thresholds . . . 52
fast-aps-switch . . . 53
frame-error . . . 54
frame-period . . . 55
frame-period-summary . . . 56
interface (OAM Link-Fault Management) . . . 57
link-adjacency-loss . . . 58
link-down . . . 58
link-discovery . . . 59
link-event-rate . . . 59
negotiation-options . . . 60
no-allow-link-events . . . 60
oam . . . 61
pdu-interval . . . 63
pdu-threshold . . . 64
protocol-down . . . 64
remote-loopback . . . 65
send-critical-event . . . 65
symbol-period . . . 66
syslog (OAM Action) . . . 66
version-ipfix (Services) . . . 67

Chapter 6 Configuration Statements: sFlow Technology 69
agent-id 69
collector 70
interfaces (sFlow Monitoring Technology) 71
polling-interval 72
sample-rate 73
source-ip 74
udp-port 74
Chapter 7 Configuration Statements: Uplink Failure Detection 75
action (Uplink Failure Detection) 75
group (Uplink Failure Detection) 76
link-to-disable 76
link-to-monitor 77
traceoptions (Uplink Failure Detection) 78
uplink-failure-detection 79
Chapter 8 Operational Commands 81
monitor traffic 82
ping 94
show oam ethernet link-fault-management 98
show interfaces (Fast Ethernet) 104
show interfaces (10-Gigabit Ethernet) 120
show sflow collector 145
show sflow 146
show sflow interface 148
traceroute 150
traceroute monitor 154
Chapter 9 Operational Commands: Analyzer 157
show forwarding-options analyzer 158

List of Figures

Part 1 Configuring Network Management

Chapter 2 **Configuring sFlow Monitoring Technology 23**

Figure 1: sFlow Technology Monitoring System 28

Chapter 3 **Configuring Uplink Failure Detection 33**

Figure 2: Uplink Failure Detection Configuration on Switches 34

List of Tables

About the Documentation xi

Table 1: Notice Icons xiii
Table 2: Text and Syntax Conventions xiii

Part 1 Configuring Network Management

Chapter 4 Routine Monitoring 37

Table 3: Output Control Keys for the monitor interface Command 39

Part 2 Configuration Statements and Operational Commands

Chapter 8 Operational Commands 81

Table 4: Match Conditions for the monitor traffic Command 84
Table 5: Logical Operators for the monitor traffic Command 86
Table 6: Arithmetic and Relational Operators for the monitor traffic Command 87
Table 7: show oam ethernet link-fault-management Output Fields 98
Table 8: show interfaces Fast Ethernet Output Fields 104
Table 9: show interfaces Gigabit Ethernet Output Fields 121
Table 10: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type 135
Table 11: show sflow collector Output Fields 145
Table 12: show sflow Output Fields 146
Table 13: show sflow interface Output Fields 148
Table 14: traceroute Output Fields 152
Table 15: traceroute monitor Output Fields 155

Chapter 9 Operational Commands: Analyzer 157

Table 16: show forwarding-options analyzer Output Fields 158

About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

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.xsl;
       }
     }
   }
   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.xsl; }
   ```

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

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

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

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

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

Documentation Conventions

Table 1 on page xiii 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 xiii 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.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** ***metric*****>;**
| (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** **(*string1* | *string2* | *string3*)**
# (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*** **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (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.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.

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

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

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

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/

- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

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

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

PART 1

Configuring Network Management

- Configuring Ethernet OAM Link Fault Management on page 3
- Configuring sFlow Monitoring Technology on page 23
- Configuring Uplink Failure Detection on page 33
- Routine Monitoring on page 37

CHAPTER 1

Configuring Ethernet OAM Link Fault Management

- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring an OAM Action Profile on page 8
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Enabling Remote Loopback Support on the Local Interface on page 15
- Configuring Link Discovery on page 16
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21

IEEE 802.3ah OAM Link-Fault Management Overview

Ethernet interfaces capable of running at 100 Mbps or faster on EX Series switches, MX Series, M Series (except M5 and M10 routers), and T Series routers support the IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology. Junos OS supports IEEE 802.3ah link-fault management.

The features of link-fault management are:

- Discovery
- Link monitoring
- Remote fault detection
- Remote loopback

The following features are not supported:

- Ethernet running on top of a Layer 2 protocol, such as Ethernet over ATM, is not supported in OAM configurations.
- Remote loopback is not supported on the 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The remote loopback feature mentioned in section 57.2.11 of IEEE 802.3ah is not supported on T4000 routers.

NOTE: **Aggregated Ethernet member links will now use the physical MAC address as the source MAC address in 802.3ah OAM packets.**

Related Documentation

- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring IEEE 802.3ah OAM Link-Fault Management

You can configure threshold values for fault events that trigger the sending of link event TLVs when the values exceed the threshold. To set threshold values for fault events on an interface, include the **event-thresholds** statement at the **[edit protocols oam ethernet link-fault-management interface]** hierarchy level.

You can also configure OAM threshold values within an action profile and apply the action profile to multiple interfaces. To create an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level.

You can configure Ethernet OAM either on an aggregate interface or on each of its member links. However, we recommend that you configure Ethernet OAM on the aggregate interface, and this will internally enable Ethernet OAM on the member links.

To view OAM statistics, use the **show oam ethernet link-fault-management** operational mode command. To clear OAM statistics, use the **clear oam ethernet link-fault-management statistics** operational mode command. To clear link-fault management state information and restart the link discovery process on Ethernet interfaces, use the **clear oam ethernet link-fault-management state** operational mode command. For more information about these commands, see the CLI Explorer.

Related Documentation

- **event-thresholds on page 52**
- *action-profile*
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15

- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling IEEE 802.3ah OAM Support

To enable IEEE 802.3ah OAM support, include the **interface** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

[edit protocols oam ethernet link-fault-management interface *interface-name*]

When you enable IEEE 802.3ah OAM on a physical interface, the discovery process is automatically triggered.

Related Documentation

- **link-fault-management on page 45**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the OAM PDU Interval

Periodic OAM PDUs are sent to perform link monitoring.

You can specify the periodic OAM PDU sending interval for fault detection.

To configure the sending interval, include the **pdu-interval** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
pdu-interval interval;
```

The periodic OAM PDU interval range is from 100 through 1000 milliseconds. The default sending interval is 1000 milliseconds.

Related Documentation

- **pdu-interval on page 63**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring the OAM PDU Threshold

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

To configure the number of PDUs that can be missed from the peer, include the **pdu-threshold** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
pdu-threshold threshold-value;
```

The threshold value range is from 3 through 10. The default is three PDUs.

Related Documentation

- pdu-threshold on page 64
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring an OAM Action Profile

You can create an action profile to define event fault flags and thresholds and the action to be taken. You can then apply the action profile to one or more interfaces.

To configure an action profile, include the **action-profile** statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level:

```
action-profile profile-name {
  action {
    syslog;
    link-down;
    send-critical-event;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
```

```
        }
        protocol-down;
    }
}
```

NOTE: Whenever link-fault management (LFM) with an action profile is configured to mark the interface as down (by including the link-down statement at the [edit protocols oam ethernet link-fault-management] hierarchy level), the port is placed in the blocked state (STP state). In such a state of the interface, data traffic is not transmitted out on that interface. Because the connectivity-fault management (CFM) downstream maintenance MEPs come up on blocked ports, the CFM sessions come up properly. However, the interface is down and the interface status TLV does not contain the correct status. Only if you configure the port status TLV, the actual status of the port is reflected. The interface status TLV does not carry the actual state of the port.

Related Documentation

- action-profile on page 47
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Threshold Values for Fault Events in an Action Profile

You can configure link event thresholds for received error events that trigger the action specified in the **action** statement. You can then apply the action profile to one or more interfaces.

To configure link event thresholds, include the **link-event-rate** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Related Documentation

- link-event-rate on page 59
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Applying an Action Profile

You can apply an action profile to one or more interfaces.

To apply an action profile to an interface, include the **apply-action-profile** statement at the **[edit protocols oam ethernet link-fault-management action-profile interface *interface-name*]** hierarchy level:

[edit protocol oam ethernet link-fault-management interface *interface-name*]
apply-action-profile *profile-name*;

Related Documentation

- **apply-action-profile on page 48**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Setting a Remote Interface into Loopback Mode

You can configure the software to set the remote DTE into loopback mode on the following interfaces:

- IQ2 and IQ2-E Gigabit Ethernet interfaces
- Ethernet interfaces on the MX Series routers or EX Series switches

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 to the management plane and processed.

To configure remote loopback, include the **remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
remote-loopback;
```

To take the remote DTE out of loopback mode, remove the **remote-loopback** statement from the configuration.

Related Documentation

- remote-loopback on page 65
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Monitoring the Loss of Link Adjacency

You can specify actions be taken when link adjacency is lost. When link adjacency is lost, the system takes the action defined in the **action** statement of the action profile.

To configure the system to take action when link adjacency is lost, include the **link-adjacency-loss** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
```

link-adjacency-loss;

Related Documentation

- link-adjacency-loss on page 58
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Monitoring Protocol Status

The CCC-DOWN flag is associated with a circuit cross-connect (CCC) connection, Layer 2 circuit, and Layer 2 VPN, which send the CCC-DOWN status to the kernel. The CCC-DOWN flag indicates that the CCC is down. The CCC-DOWN status is sent to the kernel when the CCC connection, Layer 2 circuit, or Layer 2 VPN is down. This in turn, brings down the CE-facing PE interface associated with the CCC connection, Layer 2 circuit, or Layer 2 VPN.

When the CCC-DOWN flag is signaled to the IEEE 802.3ah protocol, the system takes the action defined in the **action** statement of the action profile. For additional information about Layer 2 circuits, see the Junos OS Layer 2 Circuits Feature Guide, Junos OS VPNs Configuration Guide.

To monitor the IEEE 802.3ah protocol, on the CE-facing PE interface, include the **protocol-down** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name* event]** hierarchy level:

1. In configuration mode, go to the **[edit protocols oam ethernet link-fault-management action-profile profile-name event]** hierarchy level.

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

2. Include the **protocol-down** statement.

   ```
   [edit protocols oam ethernet link-fault-management action-profile profile-name event]
   user@host# set protocol-down
   ```

NOTE: If multiple events are specified in the action profile, all the events must occur before the specified action is taken.

Related Documentation

- protocol-down on page 64
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Enabling Remote Loopback Support on the Local Interface

You can allow a remote DTE to set a local interface into remote loopback mode on IQ2 and IQ2-E Gigabit Ethernet interfaces and all Ethernet interfaces on the MX Series routers and EX Series switches. When a remote-loopback request is sent by a remote DTE, the Junos OS places the local interface into loopback mode. When an interface is in loopback mode, all frames except OAM PDUs are looped back without any changes to the frames. OAM PDUs continue to be sent to the management plane and processed. By default, the remote loopback feature is not enabled.

To enable remote loopback, include the **allow-remote-loopback** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
   negotiation-options]
allow-remote-loopback;
```

NOTE: **Activation of OAM remote loopback may result in data frame loss.**

Related Documentation

- **allow-remote-loopback on page 48**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Link Discovery

When the IEEE 802.3ah OAM protocol is enabled on a physical interface, the discovery process is automatically triggered. 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. The discovery process is triggered automatically when OAM IEEE 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

To configure the discovery mode, include the **link-discovery** statement at the **[edit protocol oam ethernet link-fault-management interface *interface-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
link-discovery (active | passive);
```

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.

Related Documentation

- link-discovery on page 59
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Configuring Threshold Values for Local Fault Events on an Interface

You can configure threshold values on an interface for the local errors that trigger the sending of link event TLVs.

To set the error threshold values for sending event TLVs, include the **frame-error**, **frame-period**, **frame-period-summary**, and **symbol-period** statements at the **[edit protocols oam ethernet link-fault-management interface *interface-name* event-thresholds]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
event-thresholds {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Related Documentation

- event-thresholds on page 52
- frame-error on page 54
- frame-period on page 55
- frame-period-summary on page 56
- symbol-period on page 66
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15

- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Disabling the Sending of Link Event TLVs

You can disable the sending of link event TLVs.

To disable the monitoring and sending of PDUs containing link event TLVs in periodic PDUs, include the **no-allow-link-events** statement at the **[edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
   negotiation-options]
no-allow-link-events;
```

Related Documentation

- **no-allow-link-events on page 60**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Detecting Remote Faults

Fault detection is either based on flags or fault event type, length, and values (TLVs) received in OAM protocol data units (PDUs). Flags that trigger a link fault are:

- Critical Event
- Dying Gasp
- Link Fault

The link event TLVs are sent by the remote DTE by means of event notification PDUs. Link event TLVs are:

- Errored Symbol Period Event
- Errored Frame Event
- Errored Frame Period Event
- Errored Frame Seconds Summary Event

Related Documentation

- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Specifying the Actions to Be Taken for Link-Fault Management Events

You can specify the action to be taken by the system when the configured link-fault event occurs. Multiple action profiles can be applied to a single interface. For each action-profile, at least one event and one action must be specified. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all the actions are executed.

You might want to set a lower threshold for a specific action such as logging the error and set a higher threshold for another action such as sending a critical event TLV.

To specify the action, include the **action** statement at the **[edit protocols oam ethernet link-fault-management action-profile *profile-name*]** hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
event {
   link-adjacency-loss;
   protocol-down;
}
action {
   syslog;
   link-down;
   send-critical-event;
}
```

To create a system log entry when the link-fault event occurs, include the **syslog** statement.

To administratively disable the link when the link-fault event occurs, include the **link-down** statement.

To send IEEE 802.3ah link event TLVs in the OAM PDU when a link-fault event occurs, include the **send-critical-event** statement.

NOTE: If multiple actions are specified in the action profile, all of the actions are executed in no particular order.

Related Documentation

- action on page 46
- syslog on page 66
- link-down on page 58
- send-critical-event on page 65
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7
- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8

- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 21
- *Ethernet Interfaces Feature Guide for Routing Devices*

Example: Configuring IEEE 802.3ah OAM Support on an Interface

Configure 802.3ah OAM support on a 10-Gigabit Ethernet interface:

```
[edit]
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface xe-0/0/0 {
          link-discovery active;
          pdu-interval 800;
          pdu-threshold 4;
          remote-loopback;
          negotiation-options {
            allow-remote-loopback;
          }
          event-thresholds {
            frame-error 30;
            frame-period 50;
            frame-period summary 40;
            symbol-period 20;
          }
        }
      }
    }
  }
}
```

Related Documentation

- **link-fault-management on page 45**
- IEEE 802.3ah OAM Link-Fault Management Overview on page 3
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 5
- Enabling IEEE 802.3ah OAM Support on page 6
- Configuring Link Discovery on page 16
- Configuring the OAM PDU Interval on page 6
- Configuring the OAM PDU Threshold on page 7

- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Disabling the Sending of Link Event TLVs on page 18
- Detecting Remote Faults on page 18
- Configuring an OAM Action Profile on page 8
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19
- Monitoring the Loss of Link Adjacency on page 12
- Monitoring Protocol Status on page 13
- Configuring Threshold Values for Fault Events in an Action Profile on page 10
- Applying an Action Profile on page 10
- Setting a Remote Interface into Loopback Mode on page 11
- Enabling Remote Loopback Support on the Local Interface on page 15
- *Ethernet Interfaces Feature Guide for Routing Devices*

CHAPTER 2

Configuring sFlow Monitoring Technology

- Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 23
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27

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 23
- Adaptive Sampling on page 24
- sFlow Agent Address Assignment on page 25

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 27

- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26
- *Monitoring Interface Status and Traffic*

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 ip4/ipv6-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 ip4/ipv6 -address
```

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

```
[edit protocols sflow]
user@switch# set source-ip ip4/ipv6-address
```

Related Documentation

- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 23

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 28
- Overview and Topology on page 28
- Configuration on page 28
- Verification on page 30

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 1 on page 28 depicts the basic elements of the sFlow system.

Figure 1: sFlow Technology Monitoring System

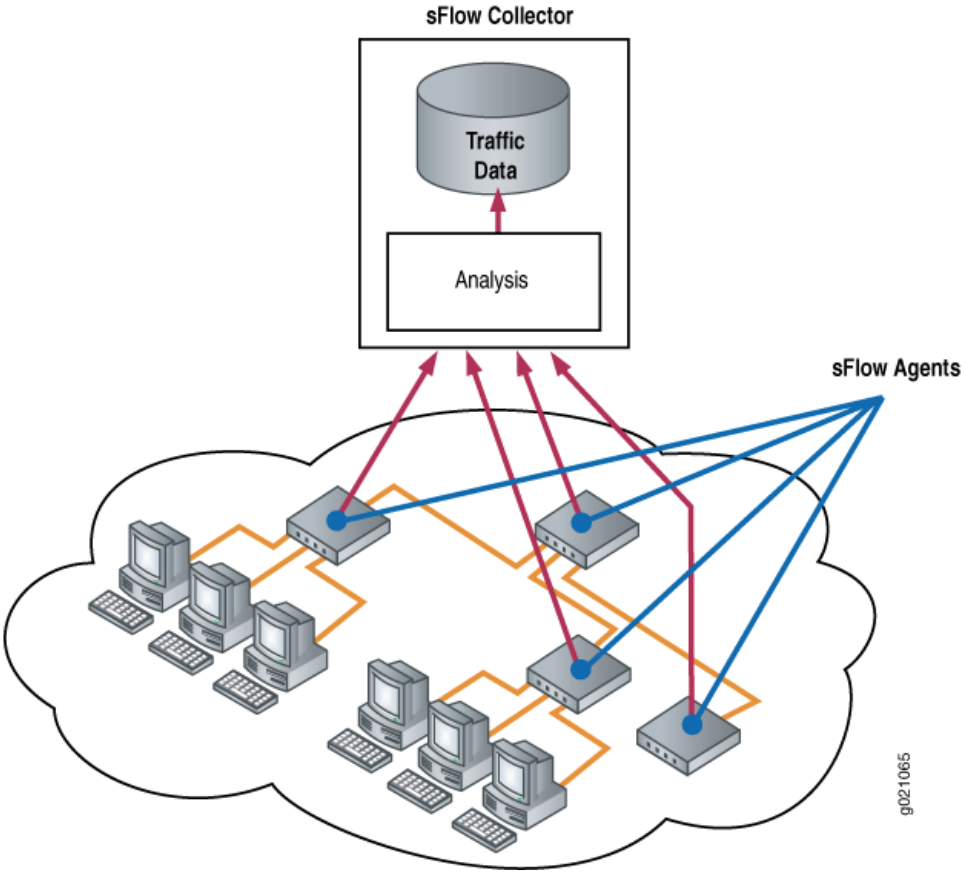

Configuration

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 30
- Verifying That sFlow Technology Is Enabled on the Specified Interface on page 31
- Verifying the sFlow Collector Configuration on page 31

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         Udp-port    No. of samples
 address
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 26
- Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch on page 23

CHAPTER 3

Configuring Uplink Failure Detection

- Understanding Uplink Failure Detection on page 33
- Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 35

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 33
- Failure Detection Pair on page 34

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 2 on page 34 illustrates a typical setup for uplink failure detection.

Figure 2: Uplink Failure Detection Configuration on Switches

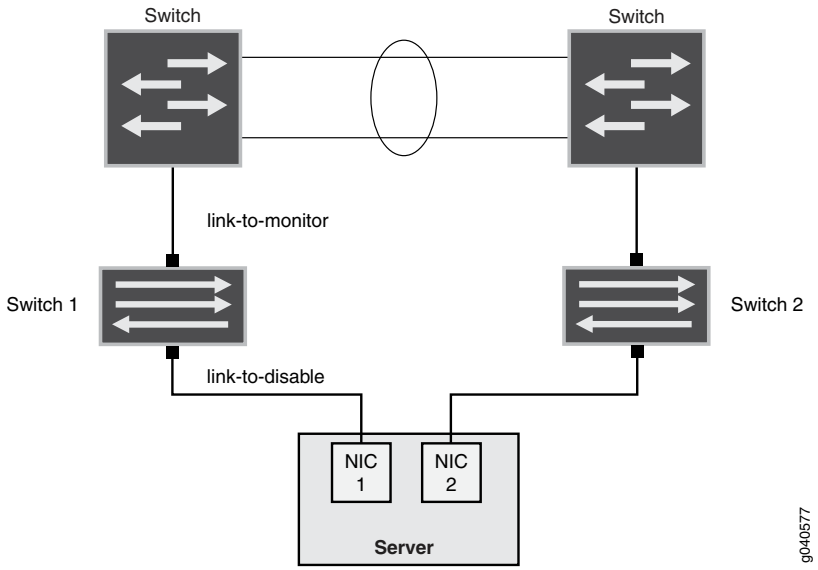

For uplink failure detection, you specify a group of uplink interfaces to be monitored and downlink interfaces to be brought down when an uplink fails. The downlink interfaces are bound to the uplink interfaces within the group. If all uplink interfaces in a group go down, then the switch brings down 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 35

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*
- Understanding Uplink Failure Detection on page 33

CHAPTER 4

Routine Monitoring

- Monitoring Traffic Through the Router or Switch on page 37

Monitoring Traffic Through the Router or Switch

To help with the diagnosis of a problem, display real-time statistics about the traffic passing through physical interfaces on the router or switch.

To display real-time statistics about physical interfaces, perform these tasks:

1. Displaying Real-Time Statistics About All Interfaces on the Router or Switch on page 37
2. Displaying Real-Time Statistics About an Interface on the Router or Switch on page 38

Displaying Real-Time Statistics About All Interfaces on the Router or Switch

Purpose Display real-time statistics about traffic passing through all interfaces on the router or switch.

Action To display real-time statistics about traffic passing through all interfaces on the router or switch:

```
user@host> monitor interface traffic
```

Sample Output

```
user@host> monitor interface traffic
host name                  Seconds: 15                 Time: 12:31:09
Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-1/0/0    Down              0          (0)                  0          (0)
 so-1/1/0    Down              0          (0)                  0          (0)
 so-1/1/1    Down              0          (0)                  0          (0)
 so-1/1/2    Down              0          (0)                  0          (0)
 so-1/1/3    Down              0          (0)                  0          (0)
 t3-1/2/0    Down              0          (0)                  0          (0)
 t3-1/2/1    Down              0          (0)                  0          (0)
 t3-1/2/2    Down              0          (0)                  0          (0)
 t3-1/2/3    Down              0          (0)                  0          (0)
 so-2/0/0      Up         211035           (1)              36778          (0)
 so-2/0/1      Up         192753           (1)              36782          (0)
 so-2/0/2      Up         211020           (1)              36779          (0)
 so-2/0/3      Up         211029           (1)              36776          (0)
 so-2/1/0      Up         189378           (1)              36349          (0)
 so-2/1/1    Down              0          (0)              18747          (0)
```

```
 so-2/1/2    Down              0          (0)              16078          (0)
 so-2/1/3      Up              0          (0)              80338          (0)
 at-2/3/0      Up              0          (0)                  0          (0)
 at-2/3/1    Down              0          (0)                  0          (0)
Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

Meaning The sample output displays traffic data for active interfaces and the amount that each field has changed since the command started or since the counters were cleared by using the **C** key. In this example, the **monitor interface** command has been running for 15 seconds since the command was issued or since the counters last returned to zero.

Displaying Real-Time Statistics About an Interface on the Router or Switch

Purpose Display real-time statistics about traffic passing through an interface on the router or switch.

Action To display traffic passing through an interface on the router or switch, use the following Junos OS CLI operational mode command:

user@host> **monitor interface** ***interface-name***

Sample Output

```
user@host> monitor interface so-0/0/1
Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
R1
Interface: so-0/0/1, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC3 Traffic statistics:
  Input bytes:                  5856541 (88 bps)
  Output bytes:                 6271468 (96 bps)
  Input packets:                 157629 (0 pps)
  Output packets:                157024 (0 pps)
Encapsulation statistics:
  Input keepalives:               42353
  Output keepalives:              42320
  LCP state: Opened
Error statistics:
  Input errors:                       0
  Input drops:                        0
  Input framing errors:               0
  Input runts:                        0
  Input giants:                       0
  Policed discards:                   0
  L3 incompletes:                     0
  L2 channel errors:                  0
  L2 mismatch timeouts:               0
  Carrier transitions:                1
  Output errors:                      0
  Output drops:                       0
  Aged packets:                       0
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                           1
  LOF count                           1
  SEF count                           1
  ES-S                               77
  SES-S                              77
```

```
SONET statistics:
  BIP-B1                                0
  BIP-B2                                0
  REI-L                                 0
  BIP-B3                                0
  REI-P                                 0
Received SONET overhead:  F1        : 0x00  J0        : 0xZ
```

Meaning The sample output shows the input and output packets for a particular SONET interface (**so-0/0/1**). The information can include common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors. For more information, see *Checklist for Tracking Error Conditions*.

To control the output of the command while it is running, use the keys shown in Table 3 on page 39.

Table 3: Output Control Keys for the monitor interface Command

Action	Key
Display information about the next interface. The **monitor interface** command scrolls through the physical or logical interfaces in the same order that they are displayed by the **show interfaces terse** command.	**N**
Display information about a different interface. The command prompts you for the name of a specific interface.	**I**
Freeze the display, halting the display of updated statistics.	**F**
Thaw the display, resuming the display of updated statistics.	**T**
Clear (zero) the current delta counters since **monitor interface** was started. It does not clear the accumulative counter.	**C**
Stop the **monitor interface** command.	**Q**

See the CLI Explorer for details on using match conditions with the **monitor traffic** command.

PART 2

Configuration Statements and Operational Commands

- Configuration Statements: Ethernet OAM Link Fault Management on page 43
- Configuration Statements: sFlow Technology on page 69
- Configuration Statements: Uplink Failure Detection on page 75
- Operational Commands on page 81
- Operational Commands: Analyzer on page 157

CHAPTER 5

Configuration Statements: Ethernet OAM Link Fault Management

- link-fault-management on page 45
- action (OAM) on page 46
- action-profile (Defining for LFM) on page 47
- allow-remote-loopback on page 48
- apply-action-profile on page 48
- ethernet (Protocols OAM) on page 49
- event (LFM) on page 52
- event-thresholds on page 52
- fast-aps-switch on page 53
- frame-error on page 54
- frame-period on page 55
- frame-period-summary on page 56
- interface (OAM Link-Fault Management) on page 57
- link-adjacency-loss on page 58
- link-down on page 58
- link-discovery on page 59
- link-event-rate on page 59
- negotiation-options on page 60
- no-allow-link-events on page 60
- oam on page 61
- pdu-interval on page 63
- pdu-threshold on page 64
- protocol-down on page 64
- remote-loopback on page 65
- send-critical-event on page 65
- symbol-period on page 66

- syslog (OAM Action) on page 66
- version-ipfix (Services) on page 67

link-fault-management

Syntax

```
link-fault-management {
  action-profile profile-name {
    action {
      link-down;
      send-critical-event;
      syslog;
    }
    event {
      link-adjacency-loss;
      link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      protocol-down;
    }
  }
  interface interface-name {
    apply-action-profile profile-name;
    link-discovery (active | passive);
    loopback-tracking;
    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 8.2.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers and EX Series switches, specify fault signaling and detection for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

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
- Enabling IEEE 802.3ah OAM Support on page 6

action (OAM)

Syntax

```
action {
  link-down;
  send-critical-event;
  syslog;
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile]

Release Information Statement introduced in Junos OS Release 8.5.

Description Define the action or actions to be taken when the OAM fault event occurs.

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

Related Documentation
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19

action-profile (Defining for LFM)

Syntax

```
action-profile profile-name {
  action {
    link-down;
    send-critical-event;
    syslog;
  }
  event {
    link-adjacency-loss;
    link-event-rate {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    protocol-down;
  }
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure a name, one or more actions, and the events that trigger the action for an action profile.

Options *profile-name*—Name of the action profile.

The remaining statements are explained separately.

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

Related Documentation

- Configuring an OAM Action Profile on page 8

allow-remote-loopback

Syntax allow-remote-loopback;

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name* negotiation-options]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable the remote loopback on IQ2 and IQ2-E Gigabit Ethernet interfaces, and Ethernet interfaces on the MX Series routers and EX Series switches.

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

Related Documentation
- Enabling Remote Loopback Support on the Local Interface on page 15

apply-action-profile

Syntax apply-action-profile *profile-name*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface]

Release Information Statement introduced in Junos OS Release 8.5.

Description Apply the specified action profile to the interface for link-fault management.

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

Related Documentation
- Applying an Action Profile on page 10

ethernet (Protocols OAM)

Syntax

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
    }
    performance-monitoring {
      delegate-server-processing;
      hardware-assisted-timestamping;
      hardware-assisted-keepalives;
      sla-iterator-profiles {
        profile-name {
          avg-fd-twoway-threshold;
          avg-ifdv-twoway-threshold;
          avg-flr-forward-threshold;
          avg-flr-backward-threshold;
          disable;
          calculation-weight {
            delay delay-weight;
            delay-variation delay-variation-weight;
          }
          cycle-time milliseconds;
          iteration-period connections;
          measurement-type (loss | statistical-frame-loss | two-way-delay);
        }
      }
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      name-format (character-string | none | dns | mac+2octet);
      maintenance-association ma-name {
        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
        protect-maintenance-association protect-ma-name;
        remote-maintenance-association remote-ma-name;
        continuity-check {
          convey-loss-threshold;
          hold-interval minutes;
          interface-status-tlv;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
          port-status-tlv;
        }
        mep mep-id {
          auto-discovery;
          direction (up | down);
          interface interface-name (protect | working);
```

```
            lowest-priority-defect (all-defects | err-xcon | mac-rem-err-xcon | no-defect |
              rem-err-xcon | xcon );
            priority number;
            remote-mep mep-id {
              action-profile profile-name;
              sla-iterator-profile profile-name {
                data-tlv-size size;
                iteration-count count-value;
                priority priority-value;
              }
            }
          }
        }
      }
    }
    evcs evc-id {
      evc-protocol cfm management-domain domain-id (management-association
        association-id | vpls (routing-instance instance-id);
      remote-uni-count count;
      multipoint-to-multipoint;
    }
    link-fault-management {
      action-profile profile-name {
        action {
          link-down;
          send-critical-event;
          syslog;
        }
        event {
          link-adjacency-loss;
          link-event-rate {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
          protocol-down;
        }
      }
      interface interface-name {
        apply-action-profile;
        link-discovery (active | passive);
        loopback-tracking;
        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;
        }
```

```
    }
  }
  lmi {
    status-counter count;
    polling-verification-timer value;
    interface name {
      uni-id uni-name;
      status-counter number;
      polling-verification-timer value;
      evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
      evc evc-name {
        default-evc;
        vlan-list vlan-id-list;
      }
    }
  }
}
```

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches, and M320, MX Series, and T Series routers, provide fault signaling and detection for 802.3ah Operation, Administration, and Management (OAM) support.

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

- Enabling IEEE 802.3ah OAM Support on page 6

event (LFM)

Syntax

```
event {
  link-adjacency-loss;
  link-event-rate {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
    }
  protocol-down;
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure threshold values for link events in an action profile.

The remaining statements are explained separately.

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

Related Documentation

- Monitoring Protocol Status on page 13

event-thresholds

Syntax

```
event-thresholds {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure threshold limit values for link events in periodic OAM PDUs.

The remaining statements are explained separately.

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 17

fast-aps-switch

Syntax fast-aps-switch;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced in Junos OS Release 12.1.

Description (M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only and EX Series switches) Reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits.

NOTE:

- Configuring this statement reduces the APS switchover time only when the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor is SAToP.
- When the **fast-aps-switch** statement is configured in revertive APS mode, you must configure an appropriate value for revert time to achieve reduction in APS switchover time.
- To prevent the logical interfaces in the data path from being shut down, configure appropriate hold-time values on all the interfaces in the data path that support TDM.
- The **fast-aps-switch** statement cannot be configured when the APS **annex-b** option is configured.
- The interfaces that have the **fast-aps-switch** statement configured cannot be used in virtual private LAN service (VPLS) environments.

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

Related Documentation

- *Reducing APS Switchover Time in Layer 2 Circuits*

frame-error

Syntax frame-error *count*;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information Statement introduced in Junos OS Release 8.4.

Description Threshold for sending frame error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value within the window.

The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.

Options *count*—Threshold count for frame error events.

Range: 0 through 100

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Configuring Threshold Values for Fault Events in an Action Profile on page 10

frame-period

Syntax frame-period *count*;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information Statement introduced in Junos OS Release 8.4.

Description Threshold for sending frame period error events or taking the action specified in the action profile.

A frame error is any frame error on the underlying physical layer. The frame period threshold is reached when the number of frame errors reaches the configured value within the period window. The default period window is the number of minimum-size frames that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

Options *count*—Threshold count for frame period error events.

Range: 0 through 100

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Configuring Threshold Values for Fault Events in an Action Profile on page 10

frame-period-summary

Syntax frame-period-summary *count*;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information Statement introduced in Junos OS Release 8.4.

Description Threshold for sending frame period summary error events or taking the action specified in the action profile.

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 window. The default window is 60 seconds. The window is not configurable.

Options *count*—Threshold count for frame period summary error events.

Range: 0 through 100

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Configuring Threshold Values for Fault Events in an Action Profile on page 10

interface (OAM Link-Fault Management)

Syntax

```
interface interface-name {
  apply-action-profile profile-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 link-fault-management]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options **interface *interface-name***—Interface to be enabled for IEEE 802.3ah link fault management OAM support.

The remaining statements are described separately.

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

Related Documentation

- Enabling IEEE 802.3ah OAM Support on page 6

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 8.5.

Description Loss of adjacency with IEEE 802.3ah link-fault management peer event. When included, the loss-of-adjacency event triggers the action specified under the **action** statement.

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

Related Documentation
- Monitoring the Loss of Link Adjacency on page 12

link-down

Syntax link-down;

Hierarchy Level [edit protocols oam ethernet link-fault-management]

Release Information Statement introduced in Junos OS Release 8.5.

Description Mark the interface down for transit traffic.

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

Related Documentation
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19

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 8.2.

Description For Ethernet interfaces on EX Series switches, and M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

Options **(active | passive)**—Passive or active mode. 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. Once the discovery process is initiated, both sides participate in discovery.

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

Related Documentation
- Configuring Link Discovery on page 16

link-event-rate

Syntax
```
link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in Junos OS Release 8.5.

Description Configure the number of link-fault management events per second.

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

Related Documentation
- Configuring Threshold Values for Fault Events in an Action Profile on page 10

negotiation-options

Syntax

```
negotiation-options {
  allow-remote-loopback;
  no-allow-link-events;
}
```

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.4.

Description Enable and disable IEEE 802.3ah Operation, Administration, and Management (OAM) features for Ethernet interfaces.

The 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

- IEEE 802.3ah OAM Link-Fault Management Overview on page 3

no-allow-link-events

Syntax no-allow-link-events;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options]

Release Information Statement introduced in Junos OS Release 8.4.

Description Disable the sending of link event TLVs.

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

Related Documentation

- Disabling the Sending of Link Event TLVs on page 18

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;
          }
```

```
        }
      }
      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;
          send-criticial-event
        }
        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]

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 802.1ag 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*
- *Example: Configuring Ethernet OAM Connectivity Fault Management on EX Series Switches*
- *Configuring Ethernet OAM Link Fault Management (CLI Procedure)*
- *Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)*

pdu-interval

Syntax pdu-interval *interval*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options ***interval***—Periodic OAM PDU sending interval.
Range: 100 through 1000 milliseconds
Default: 1000 milliseconds

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

Related Documentation

- Configuring the OAM PDU Interval on page 6

pdu-threshold

Syntax pdu-threshold *threshold-value*;

Hierarchy Level [edit protocols oam ethernet link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options ***threshold-value***—The number of PDUs missed before declaring the peer lost.

Range: 3 through 10 PDUs

Default: 3 PDUs

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

Related Documentation

- Configuring the OAM PDU Threshold on page 7

protocol-down

Syntax protocol-down;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in Junos OS Release 8.5.

Description Upper layer indication of protocol down event. When the **protocol-down** statement is included, the protocol down event triggers the action specified under the **action** statement.

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

Related Documentation

- Configuring an OAM Action Profile on page 8

remote-loopback

Syntax remote-loopback;

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in Junos OS Release 8.2.

Description For Ethernet interfaces on EX Series switches and M320, M120, MX Series, and T Series routers, set the remote DTE into loopback mode. Remove the statement from the configuration to take the remote DTE out of loopback mode. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

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

Related Documentation
- Setting a Remote Interface into Loopback Mode on page 11

send-critical-event

Syntax send-critical-event;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile action]

Release Information Statement introduced in Junos OS Release 8.5.

Description Send OAM PDUs with the critical event bit set.

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

Related Documentation
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19

symbol-period

Syntax symbol-period *count*;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event, link-event-rate], [edit protocols oam link-fault-management interface *interface-name* event-thresholds]

Release Information Statement introduced in Junos OS Release 8.4.

Description Configure the threshold for sending symbol period events or taking the action specified in the action profile.

A symbol error is any symbol code error on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period window. The default period window is the number of symbols that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.

Options ***count***—Threshold count for symbol period events.

Range: 0 through 100

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

Related Documentation

- Configuring Threshold Values for Local Fault Events on an Interface on page 17
- Configuring Threshold Values for Fault Events in an Action Profile on page 10

syslog (OAM Action)

Syntax syslog;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile action]

Release Information Statement introduced in Junos OS Release 8.5.

Description Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.

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

Related Documentation

- Specifying the Actions to Be Taken for Link-Fault Management Events on page 19

version-ipfix (Services)

Syntax

```
version-ipfix {
  template template-name {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
      flow-direction;
      vlan-id;
    }
    (ipv4-template | ipv6-template | vpls-template);
    nexthop-learning (enable |disable);
    observation-domain-id
    option-refresh-rate packets packets seconds seconds;
    options-template-id
    template-id
    template-refresh-rate packets packets seconds seconds;
  }
}
```

Hierarchy Level [edit services flow-monitoring]

Release Information Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 12.R3 for EX Series switches.
Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed.

Description Specify the IPFIX output template properties to support flow monitoring.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Inline Active flow Monitoring*
- *Configuring Flow Aggregation to Use IPFIX Flow Templates*
- *Configuring Inline Flow Monitoring on PTX Series Routers*

CHAPTER 6

Configuration Statements: sFlow Technology

- agent-id on page 69
- collector on page 70
- interfaces (sFlow Monitoring Technology) on page 71
- polling-interval on page 72
- sample-rate on page 73
- source-ip on page 74
- udp-port on page 74

agent-id

Syntax agent-id *ip4/ipv6-address*;

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. You can specify an IPv4 or IPv6 address.

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

Related Documentation
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

collector

Syntax

```
collector {
  ip4/ipv6-address;
  udp-port port-number;
}
```

Hierarchy Level [edit protocols sflow]

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

Description 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.

The remaining statements are explained separately.

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

- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

interfaces (sFlow Monitoring Technology)

Syntax

```
interfaces interface-name {
  polling-interval seconds;
  sample-rate {
    egress number;
    ingress number;
  }
}
```

Hierarchy Level [edit protocols sflow]

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

Description 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.

The remaining statements are explained separately.

Options ***interface-name***—Name of the interface on which to configure sFlow parameters.

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

Related Documentation

- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

polling-interval

Syntax polling-interval *seconds*;

Hierarchy Level [edit protocols sflow],
[edit protocols sflow interfaces *interface-name*]

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. *Polling* 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 *seconds*—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

- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27

sample-rate

Syntax

```
sample-rate {
    egress number;
    ingress number;
}
```

Hierarchy Level

[edit protocols sflow],
[edit protocols sflow interfaces *interface-name*]

Release Information

Statement introduced in Junos OS Release 9.3 for EX Series switches.
The option *number* (which immediately followed **sample-rate**) is no longer available and options **egress *number*** and **ingress *number*** 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 are disabled.

Default

By default, both egress and ingress sampling rates are disabled.

Options

egress *number*—Value for egress sampling rate.

Range: 100 through 1,073,741,823

ingress *number*—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

- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27

source-ip

Syntax source-ip *ip4/ipv6-address*;

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. You can specify an IPv4 or IPv6 address.

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

Related Documentation
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

udp-port

Syntax udp-port *port-number*;

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 ***port-number***—UDP port number for this collector.
Default: 6343

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

Related Documentation
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

CHAPTER 7

Configuration Statements: Uplink Failure Detection

- action (Uplink Failure Detection) on page 75
- group (Uplink Failure Detection) on page 76
- link-to-disable on page 76
- link-to-monitor on page 77
- traceoptions (Uplink Failure Detection) on page 78
- uplink-failure-detection on page 79

action (Uplink Failure Detection)

Syntax

```
action {
  log;
}
```

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
- Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 35

group (Uplink Failure Detection)

Syntax

```
group group-name {
  link-to-monitor interface-name;
  link-to-disable interface-name;
}
```

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 *group-name*—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

- Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 35

link-to-disable

Syntax link-to-disable *interface-name*;

Hierarchy Level [edit protocols uplink-failure-detection group *group-name*]

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 *interface-name*—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

- Configuring Interfaces for Uplink Failure Detection (CLI Procedure) on page 35

link-to-monitor

Syntax link-to-monitor *interface-name*;

Hierarchy Level [edit protocols uplink-failure-detection group *group-name*]

Release Information Statement introduced in Junos OS Release 11.1 for EX Series switches.

Description Configure the uplink interfaces to be monitored for uplink failure detection. The switch can monitor a maximum of 48 uplink interfaces in a group.

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

Options ***interface-name***—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 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 35

traceoptions (Uplink Failure Detection)

Syntax

```
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag;
}
```

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

file *filename* —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks.

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 number of trace files is reached (***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** 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.

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 everything.
- **dcd**—Trace ufdd interaction with dcd.
- **groups**—Trace uplink-failure-detection group handling.
- **interface**—Trace interface notifcation handlers of ufdd.
- **parse**—Trace configuration parsing.

no-stamp—(Optional) Do not place a timestamp on any trace file.

no-world-readable—(Optional) Restricted file access to the user who created the file.

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: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes

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 35

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 35

CHAPTER 8

Operational Commands

- monitor traffic
- ping
- show oam ethernet link-fault-management
- show interfaces (Fast Ethernet)
- show interfaces (10-Gigabit Ethernet)
- show sflow collector
- show sflow
- show sflow interface
- traceroute
- traceroute monitor

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.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX 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 4 on page 84.

Table 4: Match Conditions for the monitor traffic Command

Match Type	Condition	Description
Entity	**host** [*address* | *hostname*]	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 *address*	Matches packets with source or destination addresses containing the specified network address.
	net *address* **mask** *mask*	Matches packets containing the specified network address and subnet mask.
	port (*port-number* | *port-name*)	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** *value*	Matches packets shorter than or equal to the specified value, in bytes.
	greater *value*	Matches packets longer than or equal to the specified value, in bytes.

Table 4: 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 | igrp | 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.
	proto *ip-protocol-number*	Matches packets whose headers contain the specified IP protocol number.
	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 5 on page 86.

Table 5: 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.
&&	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 6 on page 87.

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 4 on page 84.
- 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 6: 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.

Table 6: Arithmetic and Relational Operators for the monitor traffic Command *(continued)*

Arithmetic or Relational Operator	Description
=	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 88
monitor traffic detail count on page 88
monitor traffic extensive (Absolute Sequence) on page 89
monitor traffic extensive (Relative Sequence) on page 89
monitor traffic extensive count on page 89
monitor traffic interface on page 89
monitor traffic matching on page 90
monitor traffic (TX Matrix Plus Router) on page 90
monitor traffic (QFX3500 Switch) on page 91
monitor traffic matching icmp on page 91
monitor traffic matching IP protocol number on page 92
monitor traffic matching arp on page 92
monitor traffic matching port on page 93

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 203.0.113.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0)
ack 1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0]  (ttl )
 In 203.0.113.193.179 > 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 > 203.0.113.193.179:
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 203.0.113.193.179 > 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 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.example.net.syslog > sv-log-01.example.net.syslog:
SYSLOG kernel.info, length: 57
04:11:59.862303
Out IP truncated-ip - 25 bytes missing!
summit-em0.example.net.syslog >
sv-log-02.example.net.syslog: SYSLOG kernel.info, length: 57
04:11:59.923948
In IP aj-em0.example.net.65235 >
summit-em0.example.net.telnet: .
ack 1087492766 win 33304 <nop,nop,timestamp 42366734 993490>
04:11:59.923983 Out IP truncated-ip - 232 bytes missing!
summit-em0.example.net.telnet > aj-em0.example.net.65235: P 1:241(240) ack 0 win
 33304
<nop,nop,timestamp 993590 42366734>
04:12:00.022900
In IP aj-em0.exmaple.net.65235 >
summit-em0.example.net.telnet: . ack 241 win 33304 <nop,nop,timestamp 42366834
993590>
04:12:00.141204
In IP truncated-ip - 40 bytes missing!
ipg-lnx-shell1.example.net.46182 > summit-em0.example.net.telnet: P
2950530356:2950530404(48) ack 485494987 win 63712
<nop,nop,timestamp 1308555294 987086>
04:12:00.141345
Out IP summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 1:6(5)
ack 48 win 33304
<nop,nop,timestamp 993809 1308555294>
04:12:00.141572
In IP ipg-lnx-shell1.example.net.46182 >
summit-em0.example.net.telnet: .
```

```
ack 6 win 63712
<nop,nop,timestamp 1308555294 993809>
04:12:00.141597
Out IP summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 6:10(4) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.141821
In IP ipg-lnx-shell1.example.net.46182 >
summit-em0.exmaple.net.telnet: .
ack 10 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.141837 Out IP truncated-ip - 2 bytes missing!
summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 10:20(10) ack 48 win 33304
<nop,nop,timestamp 993810 1308555294>
04:12:00.142072
In IP ipg-lnx-shell1.example.net.46182 >
summit-em0.example.net.telnet: . ack 20 win 63712
<nop,nop,timestamp 1308555294 993810>
04:12:00.142089 Out IP summit-em0.example.net.telnet >
ipg-lnx-shell1.example.net.46182: P 20:28(8) ack 48 win 33304 <nop,nop,timestamp
 993810 1308555294>
04:12:00.142321
In IP ipg-lnx-shell1.exmample.net.46182 >
summit-em0.englab.example.net.telnet: .
ack 28 win 63712 <nop,nop,timestamp 1308555294 993810>
04:12:00.142337
Out IP truncated-ip - 1 bytes missing!
summit-em0.example.net.telnet >
ipg-lnx-shell.example.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.example.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.example.net.ssh
 >
172.22.16.246.telefinder: P 132:328(196) ack 1 win 65535
...
```

monitor traffic matching icmp

```
user@host> monitor traffic matching "icmp" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

09:23:17.728737  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
322, length 40
09:23:17.728780 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
322, length 40
```

```
09:23:18.735848  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
323, length 40
09:23:18.735891 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
323, length 40
09:23:19.749732  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
324, length 40
09:23:19.749775 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
324, length 40
09:23:20.749747  In IP 172.19.10.9 > 10.10.211.93: ICMP echo request, id 1, seq
325, length 40
09:23:20.749791 Out IP 10.10.211.93 > 172.19.10.9: ICMP echo reply, id 1, seq
325, length 40
...
```

monitor traffic matching IP protocol number

```
user@host> monitor traffic matching "proto 89" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

13:06:14.700311  In IP truncated-ip - 16 bytes missing! 10.94.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:16.067010  In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:16.287566  In IP truncated-ip - 20 bytes missing! 10.94.211.142 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:20.758500  In IP truncated-ip - 16 bytes missing! 10.200.211.254 > 224.0.0
.5: OSPFv2, Hello, length 56
13:06:24.309882  In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:24.396699  In IP truncated-ip - 16 bytes missing! 10.94.211.254 > 224.0.0.
5: OSPFv2, Hello, length 56
13:06:25.067386  In IP truncated-ip - 20 bytes missing! 10.94.211.142 > 224.0.0.
5: OSPFv2, Hello, length 60
13:06:29.499988  In IP truncated-ip - 16 bytes missing! 10.200.211.254 > 224.0.0
.5: OSPFv2, Hello, length 56
13:06:32.858753  In IP truncated-ip - 20 bytes missing! 10.94.211.102 > 224.0.0.
5: OSPFv2, Hello, length 60
...
```

monitor traffic matching arp

```
user@host> monitor traffic matching “arp” no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

11:57:54.664501  In arp who-has 10.10.213.109 (00:1f:d5:f3:28:30) tell 10.10.213.31
11:57:56.828387  In arp who-has 10.10.213.233 (00:24:9d:06:77:4f) tell 10.10.213.31
11:58:01.735803  In arp who-has 10.10.213.251 (88:e0:f4:1d:41:40) tell 10.10.213.31
11:58:04.663241  In arp who-has 10.10.213.254 tell 10.94.211.170
11:58:28.488191  In arp who-has 10.10.213.149 (00:e0:91:c2:ff:8d) tell 10.10.213.31
11:58:41.858612  In arp who-has 10.10.213.148 tell 10.94.211.254
11:58:42.621533  In arp who-has 10.10.213.254 (5f:5e:ac:79:49:81) tell 10.10.213.31
11:58:44.533391  In arp who-has 10.10.213.186 tell 10.94.211.254
11:58:45.170405  In arp who-has 10.10.213.186 tell 10.94.211.254
11:58:45.770512  In arp who-has 10.10.213.186 tell 10.94.211.254
```

monitor traffic matching port

```
user@host> monitor traffic matching "port 22" no-resolve
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is OFF.
Listening on me0, capture size 96 bytes

13:14:19.108089  In IP 192.0.2.22.56714 > 10.19.300.05.22: S
2210742342:2210742342(0) win 65535 <mss 1360,nop,wscale 7,nop,nop,sackOK>
13:14:19.108165 Out IP 10.19.300.05.22 > 192.0.2.22.56714: S 23075150:23075150(0)
 ack 2210742343 win 65535 <mss 1460,nop,wscale 1,sackOK,eol>
13:14:19.136883  In IP 192.0.2.22.56714 > 10.19.300.05.22: . ack 1 win 32768
13:14:19.231364 Out IP truncated-ip - 1 bytes missing! 10.19.300.05.22 >
172.29.102.9.56714: P 1:22(21) ack 1 win 33320
13:14:19.260174  In IP truncated-ip - 10 bytes missing! 192.0.2.22.56714 >
10.94.211.93.22: P 1:31(30) ack 22 win 32767
13:14:19.284865 Out IP truncated-ip - 964 bytes missing! 10.19.300.05.22 >
172.29.102.9.56714: P 22:1006(984) ack 31 win 33320
13:14:19.314549  In IP truncated-ip - 652 bytes missing! 192.0.2.22.56714 >
10.94.211.93.22: P 31:703(672) ack 1006 win 32760
13:14:19.414135 Out IP 10.19.300.05.22 > 192.0.2.22.56714: . ack 703 win 33320
13:14:19.443858  In IP 192.0.2.22.56714 > 10.19.300.05.22: P 703:719(16) ack 1006
 win 32760
13:14:19.467379 Out IP truncated-ip - 516 bytes missing! 10.19.300.05.22 >
172.29.102.9.56714: P 1006:1542(536) ack 719 win 33320
13:14:19.734097  In IP 192.0.2.22.56714 > 10.19.300.05.22: . ack 1542 win 32768
13:14:19.843574  In IP truncated-ip - 508 bytes missing! 192.0.2.22.56714 >
10.94.211.93.22: P 719:1247(528) ack 1542 win 32768
...
```

ping

List of Syntax
Syntax on page 94
Syntax (QFX Series) on page 94

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 Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description 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.

Options ***host***—IP address or hostname of the remote system to ping.

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.

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.

detail—(Optional) Include in the output the interface on which the ping reply was received.

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.

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.

inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.

inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.

interface *source-interface*—(Optional) Interface to use to send the ping requests.

interval *seconds*—(Optional) How often to send ping requests. The range of values, in seconds, is **1** through infinity. The default value is **1**.

logical-system *logical-system-name*—(Optional) Name of logical system from which to send the ping requests.

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.

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
- *Configuring Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages*

List of Sample Output
ping hostname on page 97
ping hostname rapid on page 97
ping hostname size count on page 97

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

Sample Output

ping hostname

```
user@host> ping device1.example.com
PING device1.example.com (192.0.2.0): 56 data bytes
64 bytes from 192.0.2.0: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.0.2.0: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.0.2.0: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 1192.0.2.0: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.0.2.0: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.0.2.0: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname rapid

```
user@host> ping device1.example.com rapid
PING device1.example.com (192.0.2.0): 56 data bytes
!!!!!
--- device1.example.com 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 device1.example.com size 200 count 5
PING device1.example.com (192.0.2.0): 200 data bytes
208 bytes from 192.0.2.0: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.0.2.0: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.0.2.0: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.0.2.0: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.0.2.0: icmp_seq=4 ttl=253 time=17.898 ms

--- device1.example.com ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

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 8.2.

Description On EX Series switches and M320, M120, MX Series, T320, and T640 routers, display Operation, Administration, and Management (OAM) link fault management 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

List of Sample Output
show oam ethernet link-fault-management brief on page 102
show oam ethernet link-fault-management brief (Loopback tracking) on page 103
show oam ethernet link-fault-management detail on page 103

Output Fields Table 7 on page 98 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 7: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Indicates the status of the OAM discovery state mechanism . • **Down**—Discovery mechanism is not running.. • **Running**—Discovery mechanism is running.	All levels
Discovery state	State of the discovery mechanism. If the status of the discovery mechanism is Down then the state of discovery mechanism is Fault. However, if the status of the discovery mechanism is Running then the state can be any one of the following: • **Passive Wait** • **Active Send Local** • **Send Any** • **Send Local Remote** • **Send Local Remote Ok** • **Fault**	All levels
ISSU	Specifies that the local end is undergoing a unified in-service software upgrade (ISSU).	All levels
Peer address	Address of the OAM peer.	All levels

Table 7: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the interface. Possible values are described in the "Link Flags" section under *Common Output Fields Description*. • **Remote-Stable**—Indicates remote OAM client acknowledgment of and satisfaction with local OAM state information. **False** indicates that remote DTE either has 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. • **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 unsupported. • **Link events**—Indicates whether interpreting link events is supported or unsupported on the remote peer. • **Variable requests**—Indicates whether variable requests are supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer. • **Remote in ISSU**—Indicates that the remote end is undergoing a unified in-service software upgrade (ISSU).	All levels
Loopback Tracking	Indicates that loopback detection is enabled or disabled.	All levels
Loop Status	Indicates that a loopback issue is either found, not found, or unknown when loopback tracking is enabled.	All levels
Detect LOC	Indicates that loss-of-continuity (LOC) detection is enabled or disabled.	All levels

Table 7: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
LOC status	Indicates that a LOC issue is either found, not found, or unknown when Detect LOC is enabled. Status is unknown when LOC detection is disabled.	All levels
OAM Receive Statistics		
Information	The total number of information PDUs received.	**detail**
Event	The total number of loopback control PDUs received.	**detail**
Variable request	The total number of variable request PDUs received.	**detail**
Variable response	The total number of variable response PDUs received.	**detail**
Loopback control	The total number of loopback control PDUs received.	**detail**
Organization specific	The total number of vendor organization specific PDUs received.	**detail**
OAM Transmit Statistics		
Information	The total number of information PDUs transmitted.	**detail**
Event	The total number of event notification PDUs transmitted.	**detail**
Variable request	The total number of variable request PDUs transmitted.	**detail**
Variable response	The total number of variable response PDUs transmitted.	**detail**
Loopback control	The total number of loopback control PDUs transmitted.	**detail**
Organization specific	The total 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 since 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**

Table 7: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Total errors	The number of errored symbols that have been reported in received event TLVs since 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 since 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 reported in received event TLVs since 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 since 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 since the OAM sublayer was reset.	**detail**
OAM Transmitted Symbol Error Event Information		
Events	The number of symbol error event TLVs that have been transmitted since 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 since the OAM sublayer was reset.	**detail**

Table 7: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
OAM Current Symbol Error Event Information		
Events	The number of symbol error TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.	**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 total number of symbol errors in the period reported.	**detail**
Total errors	The number of errored symbols reported in event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.	**detail**
OAM Transmitted Frame Error Event Information		
Events	The number of errored frame event TLVs that have been transmitted since 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 since the OAM sublayer was reset.	**detail**
OAM Current Frame Error Event Information		
Events	The number of errored frame event TLVs that have been generated regardless of whether the threshold for sending event TLVs has been crossed.	**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 errored frames in the period.	**detail**
Total errors	The number of errored frames detected regardless of whether the threshold for transmitting event TLVs has been crossed.	**detail**

Sample Output

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-3/1/3
  Status: Running, Discovery state: Send Any, ISSU
  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, Remote in ISSU
```

show oam ethernet link-fault-management brief (Loopback tracking)

```
user@host> show oam ethernet link-fault-management
Interface: ge-3/1/3
  Status: Running, Discovery state: Active Send Local
  Peer address: 00:00:00:00:00:00
  Flags:0x8
  Loopback tracking: Enabled,      Loop Status: Found
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
  Interface: ge-6/1/0
    Status: Running, Discovery state: Send Any, ISSU
    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 current 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
    OAM current 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, Remote in ISSU
```

show interfaces (Fast Ethernet)

Syntax
show interfaces *interface-type*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display status information about the specified Fast Ethernet interface.

Options
interface-type—On M Series and T Series routers, the interface type is **fe-*fpc/pic/port***.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level
view

List of Sample Output
show interfaces (Fast Ethernet) on page 117
show interfaces brief (Fast Ethernet) on page 118
show interfaces detail (Fast Ethernet) on page 118
show interfaces extensive (Fast Ethernet) on page 118

Output Fields
Table 8 on page 104 lists the output fields for the **show interfaces** (Fast Ethernet) command. Output fields are listed in the approximate order in which they appear.

Table 8: show interfaces Fast Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Link-mode	Type of link connection configured for the physical interface: **Full-duplex** or **Half-duplex**	**extensive**
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(GigabitEthernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command.	**detail extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
OTN FEC statistics	The forward error correction (FEC) counters provide the following statistics:. • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 31 under the show interfaces (10-Gigabit Ethernet) command. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **PHY Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • S**ES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload (signal) label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  • **Incomplete**—Ethernet interface has the speed or link mode configured.
	  • **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other routing device manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the routing device.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive**
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	**detail extensive** none

Table 8: show interfaces Fast Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	**detail extensive**
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	**detail extensive**
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about address flag (possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

Sample Output

show interfaces (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:44 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198)
    Flags: SNMP-Traps Encapsulation: ENET2
```

```
    Protocol inet, MTU: 1500
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255
```

show interfaces brief (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 brief
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Logical interface fe-0/0/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    inet  203.0.113.1/24
```

show interfaces detail (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 detail
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:3f:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:45 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                   42                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    1                    0 pps
  Active alarms  : None
  Active defects : None
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: Is-Primary, Mac-Validate-Strict
      Mac-Validate Failures: Packets: 0, Bytes: 0
      Addresses, Flags: Is-Preferred Is-Primary
      Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,
        Generation: 136
```

show interfaces extensive (Fast Ethernet)

```
user@host> show interfaces fe-0/0/0 extensive
Physical interface: fe-0/0/0, Enabled, Physical link is Up
  Interface index: 128, SNMP ifIndex: 22, Generation: 5391
  Link-level type: Ethernet, MTU: 1514, Link-mode: Full-duplex, Speed:
  100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  CoS queues     : 4 supported, 4 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:38, Hardware address: 00:00:5e:00:53:38
  Last flapped   : 2006-01-20 14:50:58 PST (2w4d 00:46 ago)
  Statistics last cleared: Never
```

```
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                   42                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    1                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Active alarms  : None
  Active defects : None
  MAC statistics:                       Receive         Transmit
    Total octets                              0               64
    Total packets                             0                1
    Unicast packets                           0                0
    Broadcast packets                         0                1
    Multicast packets                         0                0
    CRC/Align errors                          0                0
    FIFO errors                               0                0
    MAC control frames                        0                0
    MAC pause frames                          0                0
    Oversized frames                          0
    Jabber frames                             0
    Fragment frames                           0
    VLAN tagged frames                        0
    Code violations                           0
Filter statistics:
    Input packet count                        0
    Input packet rejects                      0
    Input DA rejects                          0
    Input SA rejects                          0
    Output packet count                               1
    Output packet pad count                           0
    Output packet error count                         0
    CAM destination filters: 1, CAM source filters: 0
  Autonegotiation information:
    Negotiation status: Complete
    Link partner:
        Link partner: Full-duplex, Flow control: None, Remote fault: Ok
    Local resolution:
Packet Forwarding Engine configuration:
    Destination slot: 0
CoS information:
                Bandwidth           Buffer Priority   Limit
                              %          bps   %         usec
    0 best-effort            95    950000000  95            0      low    none
    3 network-control         5     50000000   5            0      low    none
  Logical interface fe-0/0/0.0 (Index 66) (SNMP ifIndex 198) (Generation 67)
    Flags: SNMP-Traps Encapsulation: ENET2
    Protocol inet, MTU: 1500, Generation: 105, Route table: 0
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 203.0.113/24, Local: 203.0.113.1, Broadcast: 203.0.113.255,

        Generation: 136
```

show interfaces (10-Gigabit Ethernet)

Syntax

show interfaces xe-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced in Junos OS Release 8.0.

Description

(M320, M120, MX Series, and T Series routers only) Display status information about the specified 10-Gigabit Ethernet interface.

Options

xe-*fpc/pic/port*—Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index*—(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

List of Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2) on page 135
show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode) on page 138
show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC) on page 140
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode) on page 142
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only) on page 142
show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only) on page 143

Output Fields

See Table 9 on page 121 for the output fields for the **show interfaces** (10–Gigabit Ethernet) command.

Table 9: show interfaces Gigabit Ethernet Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under *Common Output Fields Description*.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
LAN-PHY mode	10-Gigabit Ethernet interface operating in Local Area Network Physical Layer Device (LAN PHY) mode. LAN PHY allows 10-Gigabit Ethernet wide area links to use existing Ethernet applications.	All levels
WAN-PHY mode	10-Gigabit Ethernet interface operating in Wide Area Network Physical Layer Device (WAN PHY) mode. WAN PHY allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and other devices intended for SONET/SDH.	All levels
Unidirectional	Unidirectional link mode status for 10-Gigabit Ethernet interface: **Enabled** or **Disabled** for parent interface; **Rx-only** or **Tx-only** for child interfaces.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	(Gigabit Ethernet interfaces) Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under *Common Output Fields Description*.	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under *Common Output Fields Description*.	All levels

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Link flags	Information about the link. Possible values are described in the "Links Flags" section under *Common Output Fields Description*.	All levels
Wavelength	(10-Gigabit Ethernet dense wavelength-division multiplexing [DWDM] interfaces) Displays the configured wavelength, in nanometers (nm).	All levels
Frequency	(10-Gigabit Ethernet DWDM interfaces only) Displays the frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	(Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces only) Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second:timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps). The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Output Rate	Output rate in bps and pps. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Egress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for egress traffic.	**detail extensive**
Ingress account overhead	Layer 2 overhead in bytes that is accounted in the interface statistics for ingress traffic.	**detail extensive** **detail extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. The value in this field also includes the Layer 2 overhead bytes for ingress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Output bytes**—Number of bytes transmitted on the interface. The value in this field also includes the Layer 2 overhead bytes for egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface. Gigabit Ethernet and 10-Gigabit Ethernet IQ PICs count the overhead and CRC bytes. For Gigabit Ethernet IQ PICs, the input byte counts vary by interface type. For more information, see Table 9 on page 121.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored by configuring the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface. NOTE: In DPCs that are not of the enhanced type, such as DPC 40x 1GE R, DPCE 20x 1GE + 2x 10GE R, or DPCE 40x 1GE R, you might notice a discrepancy in the output of the **show interfaces** command because incoming packets might be counted in the Egress queues section of the output. This problem occurs on non-enhanced DPCs because the egress queue statistics are polled from IMQ (Inbound Message Queuing) block of the I-chip. The IMQ block does not differentiate between ingress and egress WAN traffic; as a result, the combined statistics are displayed in the egress queue counters on the Routing Engine. In a simple VPLS scenorio, if there is no MAC entry in DMAC table (by sending unidirectional traffic), traffic is flooded and the input traffic is accounted in IMQ. For bidirectional traffic (MAC entry in DMAC table), if the outgoing interface is on the same I-chip then both ingress and egress statistics are counted in a combined way. If the outgoing interface is on a different I-chip or FPC, then only egress statistics are accounted in IMQ. This behavior is expected with non-enhanced DPCs	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Ingress queues	Total number of ingress queues supported on the specified interface. Displayed on IQ2 interfaces.	**extensive**
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the routing device configuration, an alarm can ring the red or yellow alarm bell on the routing device, or turn on the red or yellow alarm LED on the craft interface. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
OTN alarms	Active OTN alarms identified on the interface.	**detail extensive**
OTN defects	OTN defects received on the interface.	**detail extensive**
OTN FEC Mode	The FECmode configured on the interface. • **efec**—Enhanced forward error correction (EFEC) is configured to defect and correct bit errors. • **gfec**—G.709 Forward error correction (GFEC) mode is configured to detect and correct bit errors. • **none**—FEC mode is not configured.	**detail extensive**
OTN Rate	OTN mode. • **fixed-stuff-bytes**—Fixed stuff bytes 11.0957 Gbps. • **no-fixed-stuff-bytes**—No fixed stuff bytes 11.0491 Gbps. • **pass-through**—Enable OTN passthrough mode. • **no-pass-through**—Do not enable OTN passthrough mode.	**detail extensive**
OTN Line Loopback	Status of the line loopback, if configured for the DWDM OTN PIC. Its value can be: **enabled** or **disabled**.	**detail extensive**
OTN FEC statistics	The forward error correction (FEC) counters for the DWDM OTN PIC. • **Corrected Errors**—The count of corrected errors in the last second. • **Corrected Error Ratio**—The corrected error ratio in the last 25 seconds. For example, 1e-7 is 1 error per 10 million bits.	**detail extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
OTN FEC alarms	OTN FEC excessive or degraded error alarms triggered on the interface. • **FEC Degrade**—OTU FEC Degrade defect. • **FEC Excessive**—OTU FEC Excessive Error defect.	**detail extensive**
OTN OC	OTN OC defects triggered on the interface. • **LOS**—OC Loss of Signal defect. • **LOF**—OC Loss of Frame defect. • **LOM**—OC Loss of Multiframe defect. • **Wavelength Lock**—OC Wavelength Lock defect.	**detail extensive**
OTN OTU	OTN OTU defects detected on the interface • **AIS**—OTN AIS alarm. • **BDI**—OTN OTU BDI alarm. • **IAE**—OTN OTU IAE alarm. • **TTIM**—OTN OTU TTIM alarm. • **SF**—OTN ODU bit error rate fault alarm. • **SD**—OTN ODU bit error rate defect alarm. • **TCA-ES**—OTN ODU ES threshold alarm. • **TCA-SES**—OTN ODU SES threshold alarm. • **TCA-UAS**—OTN ODU UAS threshold alarm. • **TCA-BBE**—OTN ODU BBE threshold alarm. • **BIP**—OTN ODU BIP threshold alarm. • **BBE**—OTN OTU BBE threshold alarm. • **ES**—OTN OTU ES threshold alarm. • **SES**—OTN OTU SES threshold alarm. • **UAS**—OTN OTU UAS threshold alarm.	**detail extensive**
Received DAPI	Destination Access Port Interface (DAPI) from which the packets were received.	**detail extensive**
Received SAPI	Source Access Port Interface (SAPI) from which the packets were received.	**detail extensive**
Transmitted DAPI	Destination Access Port Interface (DAPI) to which the packets were transmitted.	**detail extensive**
Transmitted SAPI	Source Access Port Interface (SAPI) to which the packets were transmitted.	**detail extensive**
PCS statistics	(10-Gigabit Ethernet interfaces) Displays Physical Coding Sublayer (PCS) fault conditions from the WAN PHY or the LAN PHY device. • **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode. • **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.	**detail extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem, including the following: • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. For more information, see Table 10 on page 135 • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC or a cable is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **VLAN tagged frames**—Number of frames that are VLAN tagged. The system uses the TPID of 0x8100 in the frame to determine whether a frame is tagged or not. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
OTN Received Overhead Bytes	APS/PCC0: 0x02, APS/PCC1: 0x11, APS/PCC2: 0x47, APS/PCC3: 0x58 Payload Type: 0x08	**extensive**
OTN Transmitted Overhead Bytes	APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00 Payload Type: 0x08	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem. The filtering is done by the content-addressable memory (CAM) on the PIC. The filter examines a packet's source and destination MAC addresses to determine whether the packet should enter the system or be rejected. • **Input packet count**—Number of packets received from the MAC hardware that the filter processed. • **Input packet rejects**—Number of packets that the filter rejected because of either the source MAC address or the destination MAC address. • **Input DA rejects**—Number of packets that the filter rejected because the destination MAC address of the packet is not on the accept list. It is normal for this value to increment. When it increments very quickly and no traffic is entering the routing device from the far-end system, either there is a bad ARP entry on the far-end system, or multicast routing is not on and the far-end system is sending many multicast packets to the local routing device (which the routing device is rejecting). • **Input SA rejects**—Number of packets that the filter rejected because the source MAC address of the packet is not on the accept list. The value in this field should increment only if source MAC address filtering has been enabled. If filtering is enabled, if the value increments quickly, and if the system is not receiving traffic that it should from the far-end system, it means that the user-configured source MAC addresses for this interface are incorrect. • **Output packet count**—Number of packets that the filter has given to the MAC hardware. • **Output packet pad count**—Number of packets the filter padded to the minimum Ethernet size (60 bytes) before giving the packet to the MAC hardware. Usually, padding is done only on small ARP packets, but some very small IP packets can also require padding. If this value increments rapidly, either the system is trying to find an ARP entry for a far-end system that does not exist or it is misconfigured. • **Output packet error count**—Number of packets with an indicated error that the filter was given to transmit. These packets are usually aged packets or are the result of a bandwidth problem on the FPC hardware. On a normal system, the value of this field should not increment. • **CAM destination filters, CAM source filters**—Number of entries in the CAM dedicated to destination and source MAC address filters. There can only be up to 64 source entries. If source filtering is disabled, which is the default, the values for these fields should be 0.	**extensive**
PMA PHY	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **PHY Lock**—Phase-locked loop • **PHY Light**—Loss of optical signal	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
WIS section	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET error information: • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B1**—Bit interleaved parity for SONET section overhead • **SEF**—Severely errored framing • **LOL**—Loss of light • **LOF**—Loss of frame • **ES-S**—Errored seconds (section) • S**ES-S**—Severely errored seconds (section) • **SEFS-S**—Severely errored framing seconds (section)	**extensive**
WIS line	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. State other than **OK** indicates a problem. Subfields are: • **BIP-B2**—Bit interleaved parity for SONET line overhead • **REI-L**—Remote error indication (near-end line) • **RDI-L**—Remote defect indication (near-end line) • **AIS-L**—Alarm indication signal (near-end line) • **BERR-SF**—Bit error rate fault (signal failure) • **BERR-SD**—Bit error rate defect (signal degradation) • **ES-L**—Errored seconds (near-end line) • **SES-L**—Severely errored seconds (near-end line) • **UAS-L**—Unavailable seconds (near-end line) • **ES-LFE**—Errored seconds (far-end line) • **SES-LFE**—Severely errored seconds (far-end line) • **UAS-LFE**—Unavailable seconds (far-end line)	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
WIS path	(10-Gigabit Ethernet interfaces, WAN PHY mode) Active alarms and defects, plus counts of specific SONET errors with detailed information. • **Seconds**—Number of seconds the defect has been active. • **Count**—Number of times that the defect has gone from inactive to active. • **State**—State of the error. Any state other than **OK** indicates a problem. Subfields are: • **BIP-B3**—Bit interleaved parity for SONET section overhead • **REI-P**—Remote error indication • **LOP-P**—Loss of pointer (path) • **AIS-P**—Path alarm indication signal • **RDI-P**—Path remote defect indication • **UNEQ-P**—Path unequipped • **PLM-P**—Path payload label mismatch • **ES-P**—Errored seconds (near-end STS path) • **SES-P**—Severely errored seconds (near-end STS path) • **UAS-P**—Unavailable seconds (near-end STS path) • **SES-PFE**—Severely errored seconds (far-end STS path) • **UAS-PFE**—Unavailable seconds (far-end STS path)	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation.
	• **Negotiation status:**
	  • **Incomplete**—Ethernet interface has the speed or link mode configured.
	  • **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured, or does not perform autonegotiation.
	  • **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports both **PAUSE** on receive and transmit or only **PAUSE** receive).
	  • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Received path trace, Transmitted path trace	(10-Gigabit Ethernet interfaces, WAN PHY mode) SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the routing device at the other end of the fiber. The transmitted path trace value is the message that this routing device transmits.	**extensive**
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number.	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
CoS information	Information about the CoS queue for the physical interface. • **CoS transmit queue**—Queue number and its associated user-configured forwarding class name. • **Bandwidth %**—Percentage of bandwidth allocated to the queue. • **Bandwidth bps**—Bandwidth allocated to the queue (in bps). • **Buffer %**—Percentage of buffer space allocated to the queue. • **Buffer usec**—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • **Priority**—Queue priority: **low** or **high**. • **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under *Common Output Fields Description*.	All levels

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
VLAN-Tag	Rewrite profile applied to incoming or outgoing frames on the outer (**Out**) VLAN tag or for both the outer and inner (**In**) VLAN tags. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **pop**—The outer VLAN tag of the incoming frame is removed. • **swap**—The outer VLAN tag of the incoming frame is overwritten with the user specified VLAN tag information. • **push**—An outer VLAN tag is pushed in front of the existing VLAN tag. • **push-push**—Two VLAN tags are pushed in from the incoming frame. • **swap-push**—The outer VLAN tag of the incoming frame is replaced by a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame. • **swap-swap**—Both the inner and the outer VLAN tags of the incoming frame are replaced by the user specified VLAN tag value. • **pop-swap**—The outer VLAN tag of the incoming frame is removed, and the inner VLAN tag of the incoming frame is replaced by the user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame. • **pop-pop**—Both the outer and inner VLAN tags of the incoming frame are removed.	**brief detail extensive** none
Demux:	IP demultiplexing (demux) value that appears if this interface is used as the demux underlying interface. The output is one of the following: • Source Family Inet • Destination Family Inet	**detail extensive** none
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family. Possible values are described in the "Protocol Field" section under *Common Output Fields Description*.	**detail extensive** none
MTU	Maximum transmission unit size on the logical interface.	**detail extensive** none
Maximum labels	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received and transmitted on the specified interface set. • **Input bytes, Output bytes**—Number of bytes received and transmitted on the interface set. The value in this field also includes the Layer 2 overhead bytes for ingress or egress traffic on Ethernet interfaces if you enable accounting of Layer 2 overhead at the PIC level or the logical interface level. • **Input packets, Output packets**—Number of packets received and transmitted on the interface set.	**detail extensive**
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.	**extensive**
Local statistics	Number and rate of bytes and packets destined to the routing device.	**extensive**

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Transit statistics	Number and rate of bytes and packets transiting the switch. NOTE: For Gigabit Ethernet intelligent queuing 2 (IQ2) interfaces, the logical interface egress statistics might not accurately reflect the traffic on the wire when output shaping is applied. Traffic management output shaping might drop packets after they are tallied by the **Output bytes** and **Output packets** interface counters. However, correct values display for both of these egress statistics when per-unit scheduling is enabled for the Gigabit Ethernet IQ2 physical interface, or when a single logical interface is actively using a shared scheduler.	**extensive**
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table inet.0.	**detail extensive** none
Flags	Information about protocol family flags. Possible values are described in the "Family Flags" section under *Common Output Fields Description*.	**detail extensive**
Donor interface	(Unnumbered Ethernet) Interface from which an unnumbered Ethernet interface borrows an IPv4 address.	**detail extensive** none
Preferred source address	(Unnumbered Ethernet) Secondary IPv4 address of the donor loopback interface that acts as the preferred source address for the unnumbered Ethernet interface.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	**detail extensive**
Output Filters	Names of any output filters applied to this interface. If you specify a precedence value for any filter in a dynamic profile, filter precedence values appear in parenthesis next to all interfaces.	**detail extensive**
Mac-Validate Failures	Number of MAC address validation failures for packets and bytes. This field is displayed when MAC address validation is enabled for the logical interface.	**detail extensive** none
Addresses, Flags	Information about the address flags. Possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about address flag (possible values are described in the "Addresses Flags" section under *Common Output Fields Description*.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none

Table 9: show interfaces Gigabit Ethernet Output Fields *(continued)*

Field Name	Field Description	Level of Output
Generation	Unique number for use by Juniper Networks technical support only.	**detail extensive**

For Gigabit Ethernet IQ PICs, traffic and MAC statistics output varies. Table 10 on page 135 describes the traffic and MAC statistics for two sample interfaces, each of which is sending traffic in packets of 500 bytes (including 478 bytes for the Layer 3 packet, 18 bytes for the Layer 2 VLAN traffic header, and 4 bytes for cyclic redundancy check [CRC] information). In Table 10 on page 135, the **ge-0/3/0** interface is the inbound physical interface, and the **ge-0/0/0** interface is the outbound physical interface. On both interfaces, traffic is carried on logical unit **.50** (VLAN 50).

Table 10: Gigabit Ethernet IQ PIC Traffic and MAC Statistics by Interface Type

Interface Type	Sample Command	Byte and Octet Counts Include	Comments
Inbound physical interface	**show interfaces ge-0/3/0 extensive**	Traffic statistics: Input bytes: 496 bytes per packet, representing the Layer 2 packet MAC statistics: Received octets: 500 bytes per packet, representing the Layer 2 packet + 4 bytes	The additional 4 bytes are for the CRC.
Inbound logical interface	**show interfaces ge-0/3/0.50 extensive**	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	
Outbound physical interface	**show interfaces ge-0/0/0 extensive**	Traffic statistics: Input bytes: 490 bytes per packet, representing the Layer 3 packet + 12 bytes MAC statistics: Received octets: 478 bytes per packet, representing the Layer 3 packet	For input bytes, the additional 12 bytes includes 6 bytes for the destination MAC address + 4 bytes for VLAN + 2 bytes for the Ethernet type.
Outbound logical interface	**show interfaces ge-0/0/0.50 extensive**	Traffic statistics: Input bytes: 478 bytes per packet, representing the Layer 3 packet	

Sample Output

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, IQ2)

```
user@host> show interfaces xe-5/0/0 extensive
Physical interface: xe-5/0/0, Enabled, Physical link is Up
  Interface index: 177, SNMP ifIndex: 99, Generation: 178
  Link-level type: Ethernet, MTU: 1518, LAN-PHY mode, Speed: 10Gbps, Loopback:
```

```
None, Source filtering: Enabled,
  Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 4 maximum usable queues
  Schedulers     : 1024
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:f6, Hardware address: 00:00:5e:00:53:f6
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :           6970332384                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:             81050506                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Ingress traffic statistics at Packet Forwarding Engine:
   Input  bytes  :           6970299398                    0 bps
   Input  packets:             81049992                    0 pps
   Drop   bytes  :                    0                    0 bps
   Drop   packets:                    0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0,
    L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
 FIFO errors: 0, HS link CRC errors: 0,
    MTU errors: 0, Resource errors: 0
  Ingress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort             81049992             81049992                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Egress queues: 4 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont                   0                    0                    0

  Active alarms  : None
  Active defects : None
  PCS statistics                      Seconds
    Bit errors                             0
    Errored blocks                         0
```

```
MAC statistics:                      Receive         Transmit
  Total octets                    6970332384                0
  Total packets                     81050506                0
  Unicast packets                   81050000                0
  Broadcast packets                      506                0
  Multicast packets                        0                0
  CRC/Align errors                         0                0
  FIFO errors                              0                0
  MAC control frames                       0                0
  MAC pause frames                         0                0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
Filter statistics:
  Input packet count                81050506
  Input packet rejects                   506
  Input DA rejects                         0
  Input SA rejects                         0
  Output packet count                                       0
  Output packet pad count                                   0
  Output packet error count                                 0
  CAM destination filters: 0, CAM source filters: 0
Packet Forwarding Engine configuration:
  Destination slot: 5
CoS information:
  Direction : Output
  CoS transmit queue          Bandwidth              Buffer Priority   Limit
                         %            bps     %         usec
  0 best-effort         95      950000000    95         0          low     none
  3 network-control      5       50000000     5         0          low     none

  Direction : Input
  CoS transmit queue          Bandwidth              Buffer Priority   Limit
                         %            bps     %         usec
  0 best-effort         95      950000000    95         0          low     none
  3 network-control      5       50000000     5         0          low     none

Logical interface xe-5/0/0.0 (Index 71) (SNMP ifIndex 95) (Generation 195)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.100 ]  Encapsulation: ENET2
  Egress account overhead: 100
  Ingress account overhead: 90
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                  46
   Input  packets:                   0
   Output packets:                   1
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                  46
   Input  packets:                   0
   Output packets:                   1
  Transit statistics:
   Input  bytes  :                   0                  0 bps
   Output bytes  :                   0                  0 bps
```

```
    Input  packets:                    0                    0 pps
    Output packets:                    0                    0 pps
    IPv6 transit statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
   Protocol inet, MTU: 1500, Generation: 253, Route table: 0
     Addresses, Flags: Is-Preferred Is-Primary
       Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 265
   Protocol multiservice, MTU: Unlimited, Generation: 254, Route table: 0
     Flags: None
     Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, WAN PHY Mode)

```
user@host> show interfaces xe-1/0/0 extensive
Physical interface: xe-1/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 34, Generation: 47
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Loopback: Disabled
  WAN-PHY mode
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:9d, Hardware address: 00:00:5e:00:53:9d
  Last flapped   : 2005-07-07 11:22:34 PDT (3d 12:28 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    HS Link CRC errors: 0, HS Link FIFO overflows: 0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0,
    Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0,
    Resource errors: 0
  Queue counters:       Queued packets  Transmitted packets     Dropped packets
    0 best-effort                    0                    0                   0
    1 expedited-fo                   0                    0                   0
    2 assured-forw                   0                    0                   0
    3 network-cont                   0                    0                   0
  Active alarms : LOL, LOS, LBL
  Active defects: LOL, LOS, LBL, SEF, AIS-L, AIS-P
  PCS statistics             Seconds       Count
    Bit errors                     0           0
    Errored blocks                 0           0
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
    Unicast packets                          0                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
```

```
    CRC/Align errors                          0                0
    FIFO errors                               0                0
    MAC control frames                        0                0
    MAC pause frames                          0                0
    Oversized frames                          0
    Jabber frames                             0
    Fragment frames                           0
    VLAN tagged frames                        0
    Code violations                           0
  Filter statistics:
    Input packet count                        0
    Input packet rejects                      0
    Input DA rejects                          0
    Input SA rejects                          0
    Output packet count                                        0
    Output packet pad count                                    0
    Output packet error count                                  0
    CAM destination filters: 0, CAM source filters: 0
  PMA PHY:                Seconds        Count  State
    PLL lock                    0            0  OK
    PHY light               63159            1  Light Missing
WIS section:
  BIP-B1                        0            0
  SEF                      434430       434438  Defect Active
  LOS                      434430            1  Defect Active
  LOF                      434430            1  Defect Active
  ES-S                     434430
  SES-S                    434430
  SEFS-S                   434430
WIS line:
  BIP-B2                        0            0
  REI-L                         0            0
  RDI-L                         0            0  OK
  AIS-L                    434430            1  Defect Active
  BERR-SF                       0            0  OK
  BERR-SD                       0            0  OK
  ES-L                     434430
  SES-L                    434430
  UAS-L                    434420
  ES-LFE                        0
  SES-LFE                       0
  UAS-LFE                       0
WIS path:
  BIP-B3                        0            0
  REI-P                         0            0
  LOP-P                         0            0  OK
  AIS-P                    434430            1  Defect Active
  RDI-P                         0            0  OK
  UNEQ-P                        0            0  OK
  PLM-P                         0            0  OK
  ES-P                     434430
  SES-P                    434430
  UAS-P                    434420
  ES-PFE                        0
  SES-PFE                       0
  UAS-PFE                       0
Received path trace:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ................
Transmitted path trace: orissa so-1/0/0
  6f 72 69 73 73 61 20 73 6f 2d 31 2f 30 2f 30 00   orissa so-1/0/0.
  Packet Forwarding Engine configuration:
```

```
    Destination slot: 1
  CoS information:
    CoS transmit queue          Bandwidth          Buffer     Priority   Limit
                          %           bps     %    bytes
    0 best-effort        95     950000000    95        0          low     none
    3 network-control     5      50000000     5        0          low     none
```

show interfaces extensive (10-Gigabit Ethernet, DWDM OTN PIC)

```
user@host> show interfaces ge-7/0/0 extensive
Physical interface: ge-7/0/0, Enabled, Physical link is Down
  Interface index: 143, SNMP ifIndex: 508, Generation: 208
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, BPDU Error: None,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x4000
  Link flags     : None
  Wavelength     : 1550.12 nm, Frequency: 193.40 THz
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:72, Hardware address: 00:00:5e:00:53:72
  Last flapped   : 2011-04-20 15:48:54 PDT (18:39:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                    0                    0 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                    0                    0 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
    Input  packets:                   0
    Output packets:                   0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
   Carrier transitions: 2, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                    0                    0                    0

    1 expedited-fo                   0                    0                    0

    2 assured-forw                   0                    0                    0

    3 network-cont
  Queue number:         Mapped forwarding classes
    0                   best-effort
    1                   expedited-forwarding
    2                   assured-forwarding
    3                   network-control
  Active alarms  : LINK
  Active defects : LINK
  MAC statistics:                      Receive         Transmit
    Total octets                             0                0
    Total packets                            0                0
```

```
  Unicast packets                          0               0
  Broadcast packets                        0               0
  Multicast packets                        0               0
  CRC/Align errors                         0               0
  FIFO errors                              0               0
  MAC control frames                       0               0
  MAC pause frames                         0               0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
  Total octets                             0               0
  Total packets                            0               0
  Unicast packets                          0               0
  Broadcast packets                        0               0
  Multicast packets                        0               0
  CRC/Align errors                         0               0
  FIFO errors                              0               0
  MAC control frames                       0               0
  MAC pause frames                         0               0
  Oversized frames                         0
  Jabber frames                            0
  Fragment frames                          0
  VLAN tagged frames                       0
  Code violations                          0
OTN alarms            :   None
OTN defects           :   None
OTN FEC Mode       : GFEC
OTN Rate              : Fixed Stuff Bytes 11.0957Gbps
OTN Line Loopback : Enabled
OTN FEC statistics :
  Corrected Errors                                    0
  Corrected Error Ratio (          0 sec average)     0e-0
OTN FEC alarms:          Seconds        Count  State
  FEC Degrade                  0            0  OK
  FEC Excessive                0            0  OK
OTN OC:                  Seconds        Count  State
  LOS                          2            1  OK
  LOF                      67164            2  Defect Active
  LOM                      67164           71  Defect Active
  Wavelength Lock              0            0  OK
OTN OTU:
  AIS                          0            0  OK
  BDI                      65919         4814  Defect Active
  IAE                      67158            1  Defect Active
  TTIM                         7            1  OK
  SF                       67164            2  Defect Active
  SD                       67164            3  Defect Active
  TCA-ES                       0            0  OK
  TCA-SES                      0            0  OK
  TCA-UAS                     80           40  OK
  TCA-BBE                      0            0  OK
  BIP                          0            0  OK
  BBE                          0            0  OK
  ES                           0            0  OK
  SES                          0            0  OK
  UAS                        587            0  OK
  Received DAPI:
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  Received SAPI:
```

```
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted DAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
    Transmitted SAPI:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00   ................
  OTN Received Overhead Bytes:
    APS/PCC0: 0x02, APS/PCC1: 0x42, APS/PCC2: 0xa2, APS/PCC3: 0x48
    Payload Type: 0x03
  OTN Transmitted Overhead Bytes:
    APS/PCC0: 0x00, APS/PCC1: 0x00, APS/PCC2: 0x00, APS/PCC3: 0x00
    Payload Type: 0x03
  Filter statistics:
    Input packet count                        0
    Input packet rejects                      0
    Input DA rejects                          0
    Input SA rejects                          0
    Output packet count                                       0
    Output packet pad count                                   0
    Output packet error count                                 0
    CAM destination filters: 0, CAM source filters: 0
  Packet Forwarding Engine configuration:
    Destination slot: 7
  CoS information:
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                               %            bps     %           usec
    0 best-effort             95     9500000000    95              0      low
none
    3 network-control          5      500000000     5              0      low
none
    ...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode)

```
user@host> show interfaces xe-7/0/0 extensive
Physical interface: xe-7/0/0, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 212, Generation: 174
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Enabled,
  Loopback: None, Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
...
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Transmit-Only)

```
user@host> show interfaces xe-7/0/0–tx extensive
Physical interface: xe-7/0/0-tx, Enabled, Physical link is Up
  Interface index: 176, SNMP ifIndex: 137, Generation: 177
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Tx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:19 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
```

```
   Output bytes  :      322891152287160           9627472888 bps
   Input  packets:                   0                    0 pps
   Output packets:         328809727380              1225492 pps

...

  Filter statistics:
    Output packet count           328810554250
    Output packet pad count                  0
    Output packet error count                0
...

  Logical interface xe-7/0/0-tx.0 (Index 73) (SNMP ifIndex 138) (Generation 139)

    Flags: SNMP-Traps Encapsulation: ENET2
    Egress account overhead: 100
    Ingress account overhead: 90
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :      322891152287160
     Input  packets:                   0
     Output packets:         328809727380
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :                   0                    0 bps
     Output bytes  :      322891152287160           9627472888 bps
     Input  packets:                   0                    0 pps
     Output packets:         328809727380              1225492 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 147, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 10.11.12/24, Local: 10.11.12.13, Broadcast: 10.11.12.255,
Generation: 141
    Protocol multiservice, MTU: Unlimited, Generation: 148, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

show interfaces extensive (10-Gigabit Ethernet, LAN PHY Mode, Unidirectional Mode, Receive-Only)

```
user@host> show interfaces xe-7/0/0—rx extensive
Physical interface: xe-7/0/0-rx, Enabled, Physical link is Up
  Interface index: 174, SNMP ifIndex: 118, Generation: 175
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps,
Unidirectional: Rx-Only
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
```

```
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:00:5e:00:53:83, Hardware address: 00:00:5e:00:53:83
  Last flapped   : 2007-06-01 09:08:22 PDT (3d 02:31 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :      322857456303482           9627496104 bps
   Output bytes  :                   0                    0 bps
   Input  packets:         328775413751              1225495 pps
   Output packets:                   0                    0 pps

...

  Filter statistics:
    Input packet count            328775015056
    Input packet rejects                     1
    Input DA rejects                         0

...

  Logical interface xe-7/0/0-rx.0 (Index 72) (SNMP ifIndex 120) (Generation 138)

    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :      322857456303482
     Output bytes  :                   0
     Input  packets:         328775413751
     Output packets:                   0
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :                   0
     Input  packets:                   0
     Output packets:                   0
    Transit statistics:
     Input  bytes  :      322857456303482           9627496104 bps
     Output bytes  :                   0                    0 bps
     Input  packets:         328775413751              1225495 pps
     Output packets:                   0                    0 pps
     IPv6 transit statistics:
      Input  bytes  :                   0
      Output bytes  :                   0
      Input  packets:                   0
      Output packets:                   0
    Protocol inet, MTU: 1500, Generation: 145, Route table: 0
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 192.0.2/24, Local: 192.0.2.1, Broadcast: 192.0.2.255,
Generation: 139
    Protocol multiservice, MTU: Unlimited, Generation: 146, Route table: 0
      Flags: None
      Policer: Input: __default_arp_policer__
```

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

- show sflow on page 146
- show sflow interface on page 148
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

List of Sample Output show sflow collector on page 145

Output Fields Table 11 on page 145 lists the output fields for the **show sflow collector** command. Output fields are listed in the approximate order in which they appear.

Table 11: 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
```

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
none—Display default sFlow technology configuration information.

collector—(Optional) Display standard status information about the specified sFlow collector.

interface—(Optional) Display standard status information about the specified sFlow interface.

Required Privilege Level
view

Related Documentation
- show sflow interface on page 148
- show sflow collector on page 145
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

List of Sample Output
show sflow on page 147

Output Fields
Table 12 on page 146 lists the output fields for the **show sflow** command. Output fields are listed in the approximate order in which they appear.

Table 12: 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 12: 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

- show sflow on page 146
- show sflow collector on page 145
- Example: Configuring sFlow Technology to Monitor Network Traffic on EX Series Switches on page 27
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 26

List of Sample Output show sflow interface on page 148

Output Fields Table 13 on page 148 lists the output fields for the **show sflow interface** command. Output fields are listed in the approximate order in which they appear.

Table 13: 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
```

```
                Egress Ingress   Egress Ingress   Egress Ingress
ge-0/0/0.0      Enabled  Disabled  1000    2048    1000     2048           20
```

traceroute

List of Syntax

Syntax on page 150
Syntax (QFX Series and OCX Series) on page 150

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 and OCX Series)

traceroute *host*
<as-number-lookup>
<bypass-routing>
<gateway *address*>
<inet>
<inet6>
<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.
propagate-ttl option introduced in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Support for IPv6 traceroute with **as-number-lookup** introduced with Junos OS Release 16.1.

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 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 (all | *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.

monitor*host*—(Optional) Perform this operation to display real-time monitoring information.

monitor*host*—(Optional) Perform this operation to display real-time monitoring information.

monitor*host*—(Optional) Perform this operation to display real-time monitoring information.

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-proagate-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 on page 154

List of Sample Output
traceroute on page 152
traceroute as-number-lookup host on page 153
traceroute no-resolve on page 153
traceroute propogate-ttl on page 153
traceroute (Between CE Routers, Layer 3 VPN) on page 153
traceroute (Through an MPLS LSP) on page 153

Output Fields Table 14 on page 152 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 14: 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.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
Round trip time	Average round-trip time, in milliseconds (ms).

Sample Output

traceroute

```
user@host> traceroute santacruz
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  blue23 (10.168.1.254)    2.370 ms  2.853 ms  0.367 ms
 2  red14  (10.168.255.250)  0.778 ms  2.937 ms  0.446 ms
 3  yellow (10.156.169.254)  7.737 ms  89.905 ms  0.834 ms
```

traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1  10.39.1.1 (10.39.1.1)  0.779 ms  0.728 ms  0.562 ms
 2  10.39.1.6 (10.39.1.6) [AS  32]  0.657 ms  0.611 ms  0.617 ms
 3  10.100.1.1 (10.100.1.1) [AS  10, 40, 50]  0.880 ms  0.808 ms  0.774 ms


user@host> traceroute as-number-lookup 1::1
traceroute6 to 1::1 (1::1) from 2001:b8::7, 64 hops max, 12 byte packets

user@host> traceroute 2001:b8::7 as-number-lookup
traceroute6 to 2001:b8::7 (2001:b8::7) from 2001:db8::9, 64 hops max, 12 byte
packets
 1  2001:db8::10 (2001:db8::10) [AS 18] 0.657 ms  17.319 ms  0.504 ms
 2  2001:b8::7 (2001:b8::7)  0.949 ms  0.930 ms  0.739 ms
```

traceroute no-resolve

```
user@host> traceroute santacruz no-resolve
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

traceroute propogate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
     MPLS Label=299776 CoS=0 TTL=1 S=0
     MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
     MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to host2.example.com (10.255.14.179), 30 hops max, 40
      byte packets
       1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
       2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
           MPLS Label=100006 CoS=0 TTL=1 S=1
       3  host2.example.com (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
     MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

traceroute monitor

List of Syntax
Syntax on page 154
Syntax (QFX Series) on page 154

Syntax

traceroute monitor *host*
<count *value*>
<inet | inet 6>
<interval *seconds*>
<no resolve>
<size *value*>
<source *source-address*>
<summary>

Syntax (QFX Series)

traceroute monitor *host*
<count *value*>
<inet>
<inet6>
<interval *seconds*>
<no resolve>
<size *value*>
<source *source-address*>
<summary>

Release Information

Command introduced in Junos OS Release 8.0
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display live monitoring of each hop in the route that packets take to a specified network host. Use as a debugging tool to locate points of failure in a network.

Options

host—IP address or name of remote host.

count *value*—Number of ping requests, in packets, to send in summary mode. The default value is **10**.

inet | inet6—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

interval *seconds*—(Optional) Number of seconds to wait before sending ping requests. The default value is **1**.

no resolve—(Optional) Do not attempt to display addresses symbolically.

size *value*—(Optional) Receive the specified number of bytes for each packet. The range is **0** through **65468** bytes. The default value is **64**.

source *source-address*—(Optional) Source address of the outgoing ping packets.

summary—(Optional) Generate and display a summary of live monitoring of each hop on the route that packets take to a specified network host.

Required Privilege Level network

List of Sample Output traceroute monitor on page 155

Output Fields Table 15 on page 155 describes the output fields for the **traceroute monitor** command. Output fields are listed in the approximate order in which they appear.

Table 15: traceroute monitor Output Fields

Field Name	Field Description
Host	Hostname or IP address of the router at each hop.
Loss%	Percent of packet loss. The number of ping responses divided by the number of ping requests, specified as a percentage.
Snt	Number of ping requests sent to the routerat this hop.
Last	Most recent round-trip time, in milliseconds, to the router at this hop.
Avg	Average round-trip time, in milliseconds, to the router at this hop.
Best	Shortest round-trip time, in milliseconds, to the router at this hop.
Wrst	Longest round-trip time, in milliseconds, to the router at this hop.
StDev	Standard deviation of round-trip times, in milliseconds, to the router at this hop.

Sample Output

traceroute monitor

```
user@host> traceroute monitor 10.16.0.1

                                  Loss%   Snt   Last   Avg  Best  Wrst StDev
Host
1. 10.17.41.254                   0.0%    17    0.7   1.0   0.6   5.4   1.2
2. example.net                     0.0%    17    0.6   1.0   0.6   6.6   1.4
3. example2.net                  0.0%    17    0.6   0.6   0.6   0.6   0.0
```

CHAPTER 9

Operational Commands: Analyzer

- show forwarding-options analyzer

show forwarding-options analyzer

Syntax **show forwarding-options analyzer** ***analyzer-name***

Release Information Hierarchy level **[edit forwarding-options]** introduced in Junos OS Release 13.2X50-D10 (ELS).

Description Display information about analyzers configured for mirroring.

Options ***analyzer-name***—(Optional) Displays the status of a specific analyzer on the switch.

Required Privilege Level view

Related Documentation
- *Understanding Port Mirroring and Analyzers on EX4300 Switches*

List of Sample Output show forwarding-options analyzer on page 158

Output Fields Table 16 on page 158 lists the output fields for the **show forwarding-options analyzer** command. Output fields are listed in the approximate order in which they appear.

Table 16: show forwarding-options 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 forwarding-options analyzer

```
user@switch> show forwarding-options analyzer
Analyzer name                    : employee-monitor
  Mirror rate                      : 1
  Maximum packet length            : 0
  State                            : up
  Ingress monitored interfaces     : ge-0/0/0.0
```



```
Ingress monitored interfaces     : ge-0/0/1.0
Output VLAN                      : default-switch/remote-analyzer
```